Published in final edited form as:

J Am Stat Assoc. 2015 December 1; 110(512): 1479–1490. doi:10.1080/01621459.2014.960967.

Dirichlet-Laplace priors for optimal shrinkage

Anirban Bhattacharya, Debdeep Pati, Natesh S. Pillai, and David B. Dunson

Department of Statistics, Texas A&M University, Department of Statistics, Florida State University, Department of Statistics, Harvard University, Department of Statistical Science, Duke University

Anirban Bhattacharya: anirbanb@stat.tamu.edu; Debdeep Pati: debdeep@stat.fsu.edu; Natesh S. Pillai: pillas@fas.harvard.edu; David B. Dunson: dunson@stat.duke.edu

Abstract

Penalized regression methods, such as L_1 regularization, are routinely used in high-dimensional applications, and there is a rich literature on optimality properties under sparsity assumptions. In the Bayesian paradigm, sparsity is routinely induced through two-component mixture priors having a probability mass at zero, but such priors encounter daunting computational problems in high dimensions. This has motivated continuous shrinkage priors, which can be expressed as global-local scale mixtures of Gaussians, facilitating computation. In contrast to the frequentist literature, little is known about the properties of such priors and the convergence and concentration of the corresponding posterior distribution. In this article, we propose a new class of Dirichlet–Laplace priors, which possess optimal posterior concentration and lead to efficient posterior computation. Finite sample performance of Dirichlet–Laplace priors relative to alternatives is assessed in simulated and real data examples.

Keywords

Bayesian; Convergence rate; High dimensional; Lasso; L_1 ; Penalized regression; Regularization; Shrinkage prior

1. INTRODUCTION

The overwhelming emphasis in the literature on high-dimensional data analysis has been on rapidly producing point estimates with good empirical and theoretical properties. However, in many applications, it is crucial to obtain a realistic characterization of uncertainty in estimates of parameters, functions of parameters, and predictions. Usual frequentist approaches to characterize uncertainty, such as constructing asymptotic confidence regions or using the bootstrap, can break down in high-dimensional settings. For example, in regression when the number of subjects n is equal to or larger than the number of predictors p , one cannot naively appeal to asymptotic normality and resampling from the data may not provide an adequate characterization of uncertainty.

Most penalized estimators correspond to the mode of a Bayesian posterior distribution. For example, Lasso/ L_1 regularization [28] is equivalent to maximum *a posteriori* (MAP)

estimation under a Gaussian linear regression model having a double exponential (Laplace) prior on the coefficients. Given this connection, it is natural to ask whether we can use the entire posterior distribution to provide a probabilistic measure of uncertainty. In addition to providing a characterization of uncertainty, a Bayesian perspective has distinct advantages in terms of tuning parameter choice, allowing key penalty parameters to be marginalized over the posterior distribution instead of relying on cross-validation.

From a frequentist perspective, we would like to be able to choose a default shrinkage prior that leads to similar optimality properties to those shown for L_1 penalization and other approaches. However, instead of showing that a penalized estimator obtains the minimax rate under sparsity assumptions, we would like to show that the entire posterior distribution concentrates at the optimal rate, i.e., the posterior probability assigned to a shrinking neighborhood of the true parameter value converges to one, with the neighborhood size proportional to the frequentist minimax rate.

An amazing variety of shrinkage priors have been proposed in the Bayesian literature; however with essentially no theoretical justification in the high-dimensional settings for which they were designed. [14] and [6] provided conditions on the prior for asymptotic normality of linear regression coefficients allowing the number of predictors p to increase with sample size n , with [14] requiring a very slow rate of growth and [6] assuming $p \ll n$. These results required the prior to be sufficiently flat in a neighborhood of the true parameter value, essentially ruling out shrink-age priors. [3] considered shrinkage priors in providing simple sufficient conditions for posterior consistency in linear regression where the number of variables grows slower than the sample size, though no rate of contraction was provided.

In studying posterior contraction in high-dimensions, several properties of the prior distribution are critical, including the prior concentration around sparse vectors and the implied dimensionality of the prior. Studying these properties of shrinkage priors is challenging due to the lack of exact zeros, with the prior draws being sparse in only an approximate sense. This technical hurdle has prevented any previous results on posterior concentration in high-dimensional settings for shrinkage priors. Investigating these properties is critical not just in studying frequentist optimality properties of Bayesian procedures but for Bayesians in obtaining improved insight into prior elicitation. Without such a technical handle, prior selection remains an art. Our overarching goal is to obtain theory allowing design of novel priors, which are appealing from a Bayesian perspective while having frequentist optimality properties.

2. A NEW CLASS OF SHRINKAGE PRIORS

2.1 Bayesian sparsity priors in normal means problem

For concreteness, we focus on the normal means problem ([9, 11, 20]); although the methods developed in this paper generalize trivially to high-dimensional linear and generalized linear models. In the normal means setting, one aims to estimate a n -dimensional mean¹ based on a single observation corrupted with i.i.d. standard normal noise:

$$y_i = \theta_i + \varepsilon_i, \varepsilon_i \sim N(0, 1), 1 \leq i \leq n. \quad (1)$$

Let $l_0[q; n]$ denote the subset of $\mathbb{R}^n$ given by

$$l_0[q; n] = \{\theta \in \mathbb{R}^n : \#(1 \leq j \leq n : \theta_j \neq 0) \leq q\}.$$

For a vector $x \in \mathbb{R}^r$, let $\|x\|_2$ denote its Euclidean norm. If the true mean θ_0 is q_n -sparse, i.e., $\theta_0 \in l_0[q_n; n]$, with $q_n = o(n)$, the *squared minimax rate* in estimating θ_0 in l_2 norm is $2q_n \log(n/q_n)(1 + o(1))$ [11], i.e.²

$$\inf_{\hat{\theta}} \sup_{\theta_0 \in l_0[q_n; n]} E_{\theta_0} \|\hat{\theta} - \theta_0\|_2^2 \asymp q_n \log(n/q_n). \quad (2)$$

In the above display, E_{θ_0} denotes an expectation with respect to a $N_n(\theta_0, I_n)$ density. In the presence of sparsity, one loses a logarithmic factor in the ambient dimension for not knowing the locations of the zeroes. Moreover, (2) implies that one only needs a number of replicates in the order of the sparsity to consistently estimate the mean. Appropriate thresholding/penalized estimators achieve the minimax rate (2); see [9] for detailed references.

For a subset $S \subset \{1, \dots, n\}$, let $|S|$ denote the cardinality of S and define $\theta_S = (\theta_j : j \in S)$ for a vector $\theta \in \mathbb{R}^n$. Denote $\text{supp}(\theta)$ to be the *support* of θ , the subset of $\{1, \dots, n\}$ corresponding to the non-zero entries of θ . For a vector $\theta \in \mathbb{R}^n$, a natural way to incorporate sparsity is to use point mass mixture priors:

$$\theta_j \sim (1 - \pi)\delta_0 + \pi g_\theta, j=1, \dots, n, \quad (3)$$

where $\pi = \Pr(\theta_j \neq 0)$, $\mathbb{E}|\text{supp}(\theta)| \mid \pi = n\pi$ is the prior guess on model size (sparsity level), and g_θ is an absolutely continuous density on $\mathbb{R}$. These priors are highly appealing in allowing separate control of the level of sparsity and the size of the signal coefficients. If the sparsity parameter π is estimated via empirical Bayes, the posterior median of θ is a minimax-optimal estimator [20] which can adapt to arbitrary sparsity levels as long as $q_n = o(n)$.

In a fully Bayesian framework, it is common to place a beta prior on π , leading to a beta-Bernoulli prior on the model size, which conveys an automatic multiplicity adjustment [27]. In a recent paper, [9] established that prior (3) with an appropriate beta prior on π and suitable tail conditions on g_θ leads to a minimax optimal rate of *posterior contraction*, i.e., the posterior concentrates most of its mass on a ball around θ_0 of squared radius of the order of $q_n \log(n/q_n)$:

¹Following standard practice in this literature, we use n to denote the dimensionality and it should not be confused with the sample size.

²Given sequences a_n, b_n , we denote $a_n = O(b_n)$ or $a_n \lesssim b_n$ if there exists a global constant C such that $a_n \leq Cb_n$ and $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$ as $n \rightarrow \infty$.

$$E_{\theta_0} \mathbb{P}(\|\theta - \theta_0\|_2 < Ms_n | y) \rightarrow 1, \text{ as } n \rightarrow \infty, \quad (4)$$

where $M > 0$ is a constant and $s_n^2 = q_n \log(n/q_n)$. [24] obtained consistency in model selection using point-mass mixture priors with appropriate data-driven hyperparameters.

2.2 Global-local shrinkage rules

Although point mass mixture priors are intuitively appealing and possess attractive theoretical properties, posterior sampling requires a stochastic search over an enormous space, leading to slow mixing and convergence [26]. Computational issues and consideration that many of the θ_j s may be small but not exactly zero has motivated a rich literature on continuous shrinkage priors; for some flavor refer to [3, 8, 16, 18, 25]. [26] noted that essentially all such shrinkage priors can be represented as global-local (GL) mixtures of Gaussians,

$$\theta_j \sim N(0, \psi_j \tau), \psi_j \sim f, \tau \sim g, \quad (5)$$

where τ controls global shrinkage towards the origin while the local scales $\{\psi_j\}$ allow deviations in the degree of shrinkage. If g puts sufficient mass near zero and f is appropriately chosen, GL priors in (5) can intuitively approximate (3) but through a continuous density concentrated near zero with heavy tails.

GL priors potentially have substantial computational advantages over point mass priors, since the normal scale mixture representation allows for conjugate updating of θ and ψ in a block. Moreover, a number of frequentist regularization procedures such as ridge, lasso, bridge and elastic net correspond to posterior modes under GL priors with appropriate choices of f and g . For example, one obtains a double-exponential prior corresponding to the popular L_1 or lasso penalty if f is an exponential distribution. However, many aspects of shrinkage priors are poorly understood, with the lack of exact zeroes compounding the difficulty in studying basic properties, such as prior expectation, tail bounds for the number of large signals, and prior concentration around sparse vectors. Hence, subjective Bayesians face difficulties in incorporating prior information regarding sparsity, and frequentists tend to be skeptical due to the lack of theoretical justification.

This skepticism is warranted, as it is clearly the case that reasonable seeming priors can have poor performance in high-dimensional settings. For example, choosing $\pi = 1/2$ in prior (3) leads to an exponentially small prior probability of 2^{-n} assigned to the null model, so that it becomes literally impossible to override that prior informativeness with the information in the data to pick the null model. However, with a beta prior on π , this problem can be avoided [27]. In the same vein, if one places i.i.d. $N(0, 1)$ priors on the entries of θ , then the induced prior on $\|\theta\|$ is highly concentrated around $\sqrt{n}$ leading to misleading inferences on θ almost everywhere. Although these are simple examples, similar *multiplicity problems* [27] can transpire more subtly in cases where complicated models/priors are involved and hence it is fundamentally important to understand properties of the prior and the posterior in the setting of (1).

There has been a recent awareness of these issues, motivating a basic assessment of the marginal properties of shrinkage priors for a single θ_j . Recent priors such as the horseshoe [8] and generalized double Pareto [3] are carefully formulated to obtain marginals having a high concentration around zero with heavy tails. This is well justified, but as we will see below, such marginal behavior alone is not sufficient; it is necessary to study the joint distribution of θ on $\mathbb{R}^n$. With such motivation, we propose a class of Dirichlet-kernel priors in the next subsection.

2.3 Dirichlet-kernel priors

Let ϕ_0 denote the standard normal density on $\mathbb{R}$. Also, let $\text{DE}(\tau)$ denote a zero mean double-exponential or Laplace distribution with density $f(y) = (2\tau)^{-1}e^{-|y|/\tau}$ for $y \in \mathbb{R}$. Integrating out the local scales ψ_j 's, (5) can be equivalently represented as a global scale mixture of a kernel $\mathcal{K}(\cdot)$,

$$\theta_j \stackrel{\text{i.i.d.}}{\sim} \mathcal{K}(\cdot, \tau), \tau \sim g, \quad (6)$$

where $\mathcal{K}(x) = \int \psi^{-1/2} \phi_0(x/\sqrt{\psi}) g(\psi) d\psi$ is a symmetric unimodal density on $\mathbb{R}$ and $\mathcal{K}(x, \tau) := \tau^{-1/2} \mathcal{K}(x/\sqrt{\tau})$. For example, $\psi_j \sim \text{Exp}(1/2)$ corresponds to a double-exponential kernel $\mathcal{K} \equiv \text{DE}(1)$, while $\psi_j \sim \text{IG}(1/2, 1/2)$ results in a standard Cauchy kernel $\mathcal{K} \equiv \text{Ca}(0, 1)$.

These choices lead to a kernel which is *bounded* in a neighborhood of zero. However, if one instead uses a half Cauchy prior $\psi_j^{1/2} \sim \text{Ca}_+(0, 1)$, then the resulting horseshoe kernel [7, 8] is un-bounded with a singularity at zero. This phenomenon coupled with tail robustness properties leads to excellent empirical performance of the horseshoe. However, the joint distribution of θ under a horseshoe prior is understudied and further theoretical investigation is required to understand its operating characteristics. One can imagine that it concentrates more along sparse regions of the parameter space compared to common shrinkage priors since the singularity at zero potentially allows most of the entries to be concentrated around zero with the heavy tails ensuring concentration around the relatively small number of signals.

The above class of priors rely on obtaining a suitable kernel $\mathcal{K}$ through appropriate normal scale mixtures. In this article, we offer a fundamentally different class of shrinkage priors that alleviate the requirements on the kernel, while having attractive theoretical properties. In particular, our proposed class of Dirichlet-kernel (Dk) priors replaces the single global scale τ in (6) by a vector of scales $(\phi_1\tau, \dots, \phi_n\tau)$, where $\phi = (\phi_1, \dots, \phi_n)$ is constrained to lie in the $(n-1)$ dimensional simplex $\mathcal{S}^{n-1} = \{x = (x_1, \dots, x_n)^T : x_j \geq 0, \sum_{j=1}^n x_j = 1\}$ and is assigned a $\text{Dir}(a, \dots, a)$ prior:

$$\theta_j | \phi_j, \tau \sim \mathcal{K}(\cdot, \phi_j \tau), \phi \sim \text{Dir}(a, \dots, a). \quad (7)$$

In (7), $\mathcal{K}$ is any symmetric (about zero) unimodal density with exponential or heavier tails; for computational purposes, we restrict attention to the class of kernels that can be

represented as scale mixture of normals [32]. While previous shrinkage priors obtain marginal behavior similar to the point mass mixture priors (3), our construction aims at resembling the *joint distribution* of θ under a two-component mixture prior.

We focus on the Laplace kernel from now on for concreteness, noting that all the results stated below can be generalized to other choices. The corresponding hierarchical prior given τ ,

$$\theta_j | \phi, \tau \sim \text{DE}(\phi_j \tau), \phi \sim \text{Dir}(a, \dots, a), \quad (8)$$

is referred to as a Dirichlet–Laplace prior, denoted $\theta | \tau \sim \text{DL}_a(\tau)$.

To understand the role of ϕ , we undertake a study of the marginal properties of θ_j conditional on τ , integrating out ϕ_j . The results are summarized in Proposition 2.1 below.

Proposition 2.1. *If $\theta | \tau \sim \text{DL}_a(\tau)$, then the marginal distribution of θ_j given τ is unbounded with a singularity at zero for any $a < 1$. Further, in the special case $a = 1/n$, the marginal distribution is a wrapped Gamma distribution $\text{WG}(\tau^{-1}, 1/n)$, where $\text{WG}(\lambda, \alpha)$ has a density $f(x; \lambda, \alpha) \propto |x|^{\alpha-1} e^{-\lambda|x|}$ on $\mathbb{R}$.*

Thus, marginalizing over ϕ , we obtain an unbounded kernel $\mathcal{K}$, so that the marginal density of $\theta_j | \tau$ has a singularity at 0 while retaining exponential tails. A proof of Proposition 2.1 can be found in the appendix.

The parameter τ plays a critical role in determining the tails of the marginal distribution of θ_j 's. We consider a fully Bayesian framework where τ is assigned a prior g on the positive real line and learnt from the data through the posterior. Specifically, we assume a $\text{gamma}(\lambda, 1/2)$ prior on τ with $\lambda = na$. We continue to refer to the induced prior on θ implied by the hierarchical structure,

$$\theta_j | \phi, \tau \sim \text{DE}(\phi_j \tau), \phi \sim \text{Dir}(a, \dots, a), \tau \sim \text{gamma}(na, 1/2), \quad (9)$$

as a Dirichlet–Laplace prior, denoted $\theta \sim \text{DL}_a$.

There is a frequentist literature on including a local penalty specific to each coefficient. The adaptive Lasso [31, 34] relies on empirically estimated weights that are plugged in. [23] instead sample the penalty parameters from a posterior, with a sparse point estimate obtained for each draw. These approaches do not produce a full posterior distribution but focus on sparse point estimates.

2.4 Posterior computation

The proposed class of DL priors leads to straightforward posterior computation via an efficient data augmented Gibbs sampler. The DL_a prior (9) can be equivalently represented as

$$\theta_j \sim \text{N}(0, \psi_j \phi_j^2 \tau^2), \psi_j \sim \text{Exp}(1/2), \phi \sim \text{Dir}(a, \dots, a), \tau \sim \text{gamma}(na, 1/2).$$

We detail the steps in the normal means setting noting that the algorithm is trivially modified to accommodate normal linear regression, robust regression with heavy tailed residuals, probit models, logistic regression, factor models and other hierarchical Gaussian cases. To reduce auto-correlation, we rely on marginalization and blocking as much as possible. Our sampler cycles through (i) $\theta \mid \psi, \phi, \tau, y$, (ii) $\psi \mid \phi, \tau, \theta$, (iii) $\tau \mid \phi, \theta$ and (iv) $\phi \mid \theta$. We use the fact that the joint posterior of (ψ, ϕ, τ) is conditionally independent of y given θ . Steps (ii) – (iv) together give us a draw from the conditional distribution of $(\psi, \phi, \tau) \mid \theta$, since

$$[\psi, \phi, \tau \mid \theta] = [\psi \mid \phi, \tau, \theta][\tau \mid \phi, \theta][\phi \mid \theta].$$

Steps (i) – (iii) are standard and hence not derived. Step (iv) is non-trivial and we develop an efficient sampling algorithm for jointly sampling ϕ . Usual one at a time updates of a Dirichlet vector lead to tremendously slow mixing and convergence, and hence the joint update in Theorem 2.2 is an important feature of our proposed prior; a proof can be found in the Appendix. Consider the following parametrization for the three-parameter generalized inverse Gaussian (giG) distribution: $Y \sim \text{giG}(\lambda, \rho, \chi)$ if $f(y) \propto y^{\lambda-1} e^{-0.5(\rho y + \chi/y)}$ for $y > 0$.

Theorem 2.2. *The joint posterior of $\phi \mid \theta$ has the same distribution as $(T_1/T, \dots, T_n/T)$, where T_j are independently distributed according to a $\text{giG}(a-1, 1, 2|\theta_j|)$ distribution, and $T = \sum_{j=1}^n T_j$.*

Summaries of each step are provided below.

- i. To sample $\theta \mid \psi, \phi, \tau, y$, draw θ_j independently from a $N(\mu_j, \sigma_j^2)$ distribution with

$$\sigma_j^2 = \{1 + 1/(\psi_j \phi_j^2 \tau^2)\}^{-1}, \mu_j = \{1 + 1/(\psi_j \phi_j^2 \tau^2)\}^{-1} y.$$

- ii. The conditional posterior of $\psi \mid \phi, \tau, \theta$ can be sampled efficiently in a block by independently sampling $\psi_j \mid \phi, \theta$ from an inverse-Gaussian distribution $\text{iG}(\mu_j, \lambda)$ with $\mu_j = \phi_j \tau / |\theta_j|$, $\lambda = 1$.
- iii. Sample the conditional posterior of $\tau \mid \phi, \theta$ from a $\text{giG}(\lambda - n, 1, 2 \sum_{j=1}^n |\theta_j| / \phi_j)$ distribution.
- iv. To sample $\phi \mid \theta$, draw $T_1, \dots, T_n$ independently with $T_j \sim \text{giG}(a-1, 1, 2|\theta_j|)$ and set $\phi_j = T_j/T$ with $T = \sum_{j=1}^n T_j$.

3. CONCENTRATION PROPERTIES OF DIRCHLET-LAPLACE PRIORS

In this section, we study a number of properties of the joint density of the Dirichlet-Laplace prior DL_a on $\mathbb{R}^n$ and investigate the implied rate of posterior contraction (4) in the normal means setting (1). Recall the hierarchical specification of DL_a from (9). Letting $\psi_j = \phi_j \tau$ for $j = 1, \dots, n$, a standard result (see, for example, Lemma IV.3 of [33]) implies that $\psi_j \sim$

$\text{gamma}(a, 1/2)$ independently for $j = 1, \dots, n$. Therefore, (9) can be alternatively represented as³

$$\theta_j | \psi_j \sim \text{DE}(\psi_j), \psi_j \sim \text{Ga}(a, 1/2). \quad (10)$$

The formulation (10) is analytically convenient since the joint distribution factors as a product of marginals and the marginal density can be obtained analytically in Proposition 3.1 below. The proof follows from standard properties of the modified Bessel function [15]; a proof is sketched in the Appendix.

Proposition 3.1. *The marginal density Π of θ_j for any $1 \leq j \leq n$ is given by*

$$\Pi(\theta_j) = \frac{1}{2^{(1+a)/2} \Gamma(a)} |\theta_j|^{(a-1)/2} K_{1-a} \left(\sqrt{2|\theta_j|} \right), \quad (11)$$

where

$$K_\nu(x) = \frac{\Gamma(\nu+1/2)(2x)^\nu}{\sqrt{\pi}} \int_0^\infty \frac{\cos t}{(t^2+x^2)^{\nu+1/2}} dt$$

is the modified Bessel function of the second kind.

Figure 3 plots the marginal density (11) to compare with common shrinkage priors.

We continue to denote the joint density of θ on $\mathbb{R}^n$ by Π , so that $\Pi(\theta) = \prod_{j=1}^n \Pi(\theta_j)$. For a subset $S \subset \{1, \dots, n\}$, let Π_S denote the marginal distribution of $\theta_S = \{\theta_j : j \in S\} \in \mathbb{R}^{|S|}$. For a Borel set $A \subset \mathbb{R}^n$, let $\mathbb{P}(A) = \int_A \Pi(\theta) d\theta$ denote the prior probability of A , and $\mathbb{P}(A | y^{(n)})$ the posterior probability given data $y^{(n)} = (y_1, \dots, y_n)$ and the model (1). Finally, let $E_{\theta_0}/P_{\theta_0}$ respectively indicate an expectation/probability w.r.t. the $N_n(\theta_0, I_n)$ density. We now establish that under mild restrictions on $\|\theta_0\|$, the posterior arising from the DL_a prior (9) contracts at the minimax rate of convergence for an appropriate choice of the Dirichlet concentration parameter a .

Theorem 3.1. *Consider model (1) with $\theta \sim \text{DL}_{a_n}$ as in (9), where $a_n = n^{-(1+\beta)}$ for some $\beta > 0$ small. Assume $\theta_0 \in l_0[q_n; n]$ with $q_n = o(n)$ and $\|\theta_0\|_2^2 \leq q_n \log^4 n$. Then, with $s_n^2 = q_n \log(n/q_n)$ and for some constant $M > 0$,*

$$\lim_{n \rightarrow \infty} E_{\theta_0} \mathbb{P}(\|\theta - \theta_0\|_2 < M s_n | y) = 1. \quad (12)$$

If $a_n = 1/n$ instead, then (12) holds when $q_n \gtrsim \log n$.

A proof of Theorem 3.1 can be found in Section 6. Theorem 3.1 is the first result obtaining posterior contraction rates for a continuous shrinkage prior in the normal means setting or the closely related high-dimensional regression problem. Theorem 3.1 posits that when the

³This formulation only holds when $\tau \sim \text{gamma}(na, 1/2)$ and is not true for the general $\text{DL}_a(\tau)$ class with $\tau \sim g$.

parameter a in the Dirichlet–Laplace prior is chosen, depending on the sample size, to be $n^{-(1+\beta)}$ for any $\beta > 0$ small, the resulting posterior contracts at the minimax rate (2), provided $\|\theta_0\|_2^2 \leq q_n \log^4 n$. Using the Cauchy–Schwarz inequality, $\|\theta_0\|_1^2 \leq q_n \|\theta_0\|_2^2$ for $\theta_0 \in l_0[q_n; n]$ and the bound on $\|\theta_0\|_2$ implies that $\|\theta_0\|_1 \leq q_n (\log n)^2$. Hence, the condition in Theorem 3.1 permits each non-zero signal to grow at a $(\log n)^2$ rate, which is a fairly mild assumption. Moreover, in a recent technical report, the authors showed that a large subclass of global-local priors (5) including the Bayesian lasso lead to a sub-optimal rate of posterior convergence; i.e., the expression in (12) converges to 0 whenever $\|\theta_0\|_2^2/q_n \rightarrow \infty$. Therefore, Theorem 3.1 indeed provides a substantial improvement over a large class of GL priors.

The choice $a_n = n^{-(1+\beta)}$ will be evident from the various auxiliary results in Section 3.1, specifically Lemma 3.3 and Theorem 3.4. The conclusion of Theorem 3.1 continues to hold when $a_n = 1/n$ under an additional mild assumption on the sparsity q_n . In Tables 1 & 2 of Section 4, detailed empirical results are provided with $a_n = 1/n$ as a default choice. Based on empirical experience, we find that $a = 1/n$ may have a tendency to over-shrink the signals in cases where there are several relatively small signals. In such settings, depending on the practitioner’s utility function, the singularity at zero can be softened using a DL_a prior for a larger value of a . We report the results for $a = 1/2$ in Section 4, whence computational gains arise as the distribution of T_j in (iv) turns out to be inverse-Gaussian (iG), for which exact samplers are available. A fully Bayesian procedure with a discrete hyper prior on a is considered in the real data application.

In the regression setting $y \sim N_n(X\beta, I_n)$ with the number of covariates $p_n \rightarrow n$, it is typically assumed that the eigenvalues of $(X^T X)/n$ are bounded between two absolute constants [4, 19]. In such cases, Theorem 3.1 can be used to derive posterior convergence rates for the vector of regression coefficients β in l_2 norm. [4] studied posterior consistency in this setting for a class of shrinkage priors but did not provide contraction rates. While our methods can be extended trivially to the ultra high-dimensional regression setting, substantial work will be required to establish theoretical properties under standard restricted isometry type conditions on the design matrix commonly used to prove oracle inequalities for frequentist penalized methods [29]. Even with point mass mixture priors, such contraction results are unknown with the exception of a recent technical report [10]. While we propose a heuristic variable selection method which seems to work well in practice, obtaining model selection consistency results similar to [19, 24] is challenging and will be addressed elsewhere.

3.1 Auxiliary results

In this section, we state a number of properties of the DL prior which provide a better understanding of the joint prior structure and also crucially help us in proving Theorem 3.1. We first provide useful bounds on the joint density of the DL prior in Lemma 3.2 below; a proof can be found in the Appendix.

Lemma 3.2. *Consider the DL_a prior on $\mathbb{R}^n$ for a small. Let $S \subset \{1, \dots, n\}$ and $\eta \in \mathbb{R}^{|S|}$.*

If $\min_{j \in |S|} |\eta_j| > \delta$ for δ small, then

$$\log \Pi_S(\eta) \leq C|S|\log(1/\delta), \quad (13)$$

where $C > 0$ is an absolute constant.

If $\|\eta\|_2 \leq m$ for m large, then

$$-\log \Pi_S(\eta) \leq C\{|S|\log(1/a) + |S|^{3/4}m^{1/2}\}, \quad (14)$$

where $C > 0$ is an absolute constant.

It is evident from Figure 3 that the univariate marginal density Π has an infinite spike near zero. We quantify the probability assigned to a small δ -neighborhood of the origin in Lemma 3.3 below.

Lemma 3.3. Assume $\theta_1 \in \mathbb{R}$ has a probability density Π as in (11). Then, for $\delta > 0$ small,

$$\mathbb{P}(|\theta_1| > \delta) \leq C\log(1/\delta)/\Gamma(a),$$

where $C > 0$ is an absolute constant.

A proof of Lemma 3.3 can be found in the Appendix.

In case of point mass mixture priors (3), the induced prior on the model size $|\text{supp}(\theta)|$ follows a Binomial (n, π) prior (given π), facilitating study of the multiplicity phenomenon [27]. However, $\mathbb{P}(\theta = 0) = 0$ for any continuous shrinkage prior, which compounds the difficulty in studying the degree of shrinkage for these classes of priors. Letting $\text{supp}_\delta(\theta) = \{j : |\theta_j| > \delta\}$ to be the entries in θ larger than δ in magnitude, we propose $|\text{supp}_\delta(\theta)|$ as an approximate measure of model size for continuous shrinkage priors. We show in Theorem 3.4 below that for an appropriate choice of δ , $|\text{supp}_\delta(\theta)|$ doesn't exceed a constant multiple of the true sparsity level q_n with *posterior probability* tending to one, a property which we refer to as *posterior compressibility*.

Theorem 3.4. Consider model (1) with $\theta \sim \text{DL}_{a_n}$ as in (9), where $a_n = n^{-(1+\beta)}$ for some $\beta > 0$ small. Assume $\theta_0 \in l_0[q_n; n]$ with $q_n = o(n)$. Let $\delta_n = q_n/n$. Then,

$$\lim_{n \rightarrow \infty} E_{\theta_0} \mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n | y^{(n)}) = 0, \quad (15)$$

for some constant $A > 0$. If $a_n = 1/n$ instead, then (15) holds when $q_n \gtrsim \log n$.

The choice of δ_n in Theorem 3.4 guarantees that the entries in θ smaller than δ_n in magnitude produce a negligible contribution to $\|\theta\|$. Observe that the prior distribution of $|\text{supp}_{\delta_n}(\theta)|$ is Binomial (n, ζ_n) , where $\zeta_n = \mathbb{P}(|\theta_1| > \delta_n)$. When $a_n = n^{-(1+\beta)}$, ζ_n can be bounded above by $\log n/n^{1+\beta}$ in view of Lemma 3.3 and the fact that $\Gamma(x) \sim 1/x$ for x small. Therefore, the prior expectation $\mathbb{E}|\text{supp}_{\delta_n}(\theta)| \sim \log n/n^\beta$. This implies an exponential tail bound for $\mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n)$ by Chernoff's method, which is instrumental in deriving Theorem 3.4. A proof of Theorem 3.4 along these lines can be found in the Appendix.

The posterior compressibility property in Theorem 3.4 ensures that the dimensionality of the posterior distribution of θ (in an approximate sense) doesn't substantially overshoot the true dimensionality of θ_0 , which together with the bounds on the joint prior density near zero and infinity in Lemma 3.2 delivers the minimax rate in Theorem 3.1.

4. SIMULATION STUDY

To illustrate the finite-sample performance of the proposed DL prior, we tabulate the results from a replicated simulation study with various dimensionality n and sparsity level q_n in Tables 1 and 2. In each setting, we sample 100 replicates of a n -dimensional vector y from a $N_n(\theta_0, I_n)$ distribution with θ_0 having q_n non-zero entries which are all set to be a constant $A > 0$. We chose two values of n , namely $n = 100, 200$. For each n , we let the model size q_n to be 5, 10, 20% of n and vary A over 5, 6, 7, 8. This results in 24 simulation settings in total; for each setting the squared error loss corresponding to the posterior median averaged across simulation replicates is tabulated. The simulations were designed to mimic the setting in Section 3 where θ_0 is sparse with a few moderate-sized coefficients.

Based on the discussion in Section 3, we present the results for the DL prior with $a = 1/n$ and $a = 1/2$. To offer grounds for comparison, we have tabulated the results for Lasso (LS), Empirical Bayes median (EBMed)⁴ [20], posterior median with a point mass prior (PM) [9], posterior median corresponding to the Bayesian lasso [25] and the horseshoe [8]. For the fully Bayesian analysis using point mass mixture priors, we use a complexity prior on the subset-size, $\pi_n(s) \propto \exp\{-\kappa s \log(2n/s)\}$ with $\kappa = 0.1$ and independent standard Laplace priors for the non-zero entries as in [9].⁵ A wide difference between most of the competitors and the proposed $DL_{1/n}$ is observed in Table 2. As evident from Tables 1 and 2, $DL_{1/2}$ is robust to smaller signal strength as is the horseshoe.

We also illustrate a high-dimensional simulation setting akin to an example in [8], where one has a single observation y from a $n = 1000$ dimensional $N_n(\theta_0, I_n)$ distribution, with $\theta_0[1 : 10] = 10$, $\theta_0[11 : 100] = A$, and $\theta_0[101 : 1000] = 0$. We then vary A from 2 to 7 and summarize the squared error averaged across 100 replicates in Table 3. We only compare the Bayesian shrinkage priors here; the squared error for the posterior median is tabulated. Table 3 clearly illustrates the need for prior elicitation in high dimensions.

For visual illustration and comparison, we present the results from a single replicate in the first simulation setting with $n = 200$, $q_n = 10$ and $A = 7$ in Figure 2 & 3. The blue circles indicate the entries of y , while the red circles correspond to the posterior median of θ . The shaded region corresponds to a 95% point wise credible interval for θ .

5. PROSTATE DATA APPLICATION

We consider a popular dataset [12, 13] from a microarray experiment consisting of expression levels for 6033 genes for 50 normal control subjects and 52 patients diagnosed

⁴The EBMed procedure was implemented using the package [21].

⁵Given a draw for s , a subset S of size s is drawn uniformly. Set $\theta_j = 0$ for all $j \notin S$ and draw $\theta_j, j \in S$ i.i.d. from standard Laplace. The beta-bernoulli priors in (3) induce a similar prior on the subset size.

with prostate cancer. The data take the form of a 6033×102 matrix with the (i, j) th entry corresponding to the expression level for gene i on patient j ; the first 50 columns correspond to the normal control subjects with the remaining 52 for the cancer patients. The goal of the study is to discover genes whose expression levels *differ* between the prostate cancer patients (treatment) and normal subjects (control). A two sample t -test with 100 degrees of freedom was implemented for each gene and the resulting t -statistic t_i was converted to a z -statistic $z_i = \Phi^{-1}(T_{100}(t_i))$. Under the null hypothesis H_{0i} of no difference in expression levels between the treatment and control group for the i th gene, the null distribution of z_i is $N(0, 1)$. Figure 4 shows a histogram of the z -values, comparing it to a $N(0, 1)$ density with a multiplier chosen to make the curve integrate to the same area as the histogram. The shape of the histogram suggests the presence of certain interesting genes [12].

The classical Bonferroni correction for multiple testing flags only 6 genes as significant, while the two-group empirical Bayes method of [20] found 139 significant genes, being much less conservative. The local Bayes false discovery rate (fdr) [5] control method identified 54 genes as non-null. For detailed analysis of this dataset using existing methods, refer to [12, 13].

To apply our method, we set up a normal means model $z_i = \theta_i + \varepsilon_i$, $i = 1, \dots, 6033$ and assign θ a DL_a prior. Instead of fixing a , we place a discrete uniform prior on a supported on the interval $[1/6000, 1/2]$, with the support points of the form $10(k+1)/6000$, $k = 0, 1, \dots, K$. Such a fully Bayesian approach allows the data to dictate the choice of the tuning parameter a which is only specified up to a constant by the theory and also avoids potential numerical issues arising from fixing $a = 1/n$ when n is large. Updating a is straightforward since the full conditional distribution of a is again a discrete distribution on the chosen support points.

A referee pointed out that the z -values for 6033 genes obtained using $n = 102$ observations are unlikely to be independent and recommended investigating robustness of our approach under correlated errors. We conducted a small simulation study to this effect to evaluate the performance of our method when the error distribution is misspecified. We generated data from the model $y_i = \theta_i + \varepsilon_i$ with $(\varepsilon_1, \dots, \varepsilon_n)^T \sim N_n(0, \Omega)$, where Ω corresponds to the covariance matrix of an auto-regression sequence with pure error variance σ^2 and auto-regressive coefficient ρ . We placed a discrete uniform prior on a supported on the interval $[1/1000, 1/2]$, with the support points of the form $10(k+1)/1000$, $k = 0, 1, \dots, K$. We observe that the mean squared error of the posterior median for 50 replicated datasets with $n = 1000$, $q_n = 50$, $A = 4, 5, 6$, $\sigma^2 = 1$ increases by only 3% on average when ρ increases from 0 to 0.1. Therefore, the proposed method is robust to mild misspecification in the covariance structure. However, if further prior information is available regarding the covariance structure, that should be incorporated in the modeling framework.

For the real data application, we implemented the Gibbs sampler in Section 2.4 for 10,000 draws discarding a burn-in of 5,000. Mixing and convergence of the Gibbs sampler was satisfactory based on examination of trace plots, with the 5,000 retained samples having an effective sample size of 2369.2 averaged across the θ_i 's. The computational time per iteration scaled approximately linearly with the dimension. The posterior mode of a was at $1/20$.

In this application, we expect there to be two clusters of $|\theta_j|$ s, with one concentrated closely near zero corresponding to genes that are effectively not differentially expressed and another away from zero corresponding to interesting genes for further study. As a simple automated approach, we cluster $|\theta_j|$ s at each MCMC iteration using *kmeans* with 2 clusters. For each iteration, the number of non-zero signals is then estimated by the smaller cluster size out of the two clusters. A final estimate (M) of the number of non-zero signals is obtained by taking the mode over all the MCMC iterations. The M largest (in absolute magnitude) entries of the posterior median are identified as the non-zero signals.

Using the above selection scheme, our method declared 128 genes as non-null. Interestingly, out of the 128 genes, 100 are common with the ones selected by EBMed. Also all the 54 genes obtained using FDR control form a subset of the selected 128 genes. Horseshoe is overly conservative; it selected only 1 gene (index: 610) using the same clustering procedure; the selected gene was the one with the largest *effect size* (refer to Table 11.2 in [13]).

6. PROOF OF THEOREM 3.1

We prove Theorem 3.1 for $a_n = 1/n$ in details and note the places where the proof differs in case of $a_n = n^{-(1+\beta)}$. Recall $\theta_S := \{\theta_j, j \in S\}$ and for $\delta > 0$, $\text{supp}_\delta(\theta) := \{j : |\theta_j| > \delta\}$. Let $E_{\theta_0}/P_{\theta_0}$ respectively indicate an expectation/probability w.r.t. the $N_n(\theta_0, I_n)$ density.

For a sequence of positive real numbers r_n to be chosen later, let $\delta_n = r_n/n$. Define

$$\mathcal{D}_n = \int \prod_{i=1}^n f_{\theta_i}(y_i) / f_{\theta_0} d\Pi(\theta). \text{ Let}$$

$$\mathcal{A}_n = \{\mathcal{D}_n \geq e^{-4r_n^2} \mathbb{P}(\|\theta - \theta_0\|_2 \leq 2r_n)\}$$

be a subset of $\sigma(y^{(n)})$, the sigma-field generated by $y^{(n)}$, as in Lemma 5.2 of [9] such that

$P_{\theta_0}(\mathcal{A}_n^c) \leq e^{-r_n^2}$. Let $\mathcal{S}_n$ be the collection of subsets $S \subset \{1, 2, \dots, n\}$ such that $|S| \leq Aq_n$. For each such S and a positive integer j , let $\{\theta^{S,j,i} : i = 1, \dots, N_{S,j}\}$ be a $2jr_n$ net of $\Theta_{S,j,n} = \{\theta \in \mathbb{R}^n : \text{supp}_{\delta_n}(\theta) = S, 2jr_n \leq \|\theta - \theta_0\|_2 \leq 2(j+1)r_n\}$ created as follows. Let $\{\phi^{S,j,i} : i = 1, \dots, N_{S,j}\}$ be a jr_n net of the $|S|$ -dimensional ball $\{\|\phi - \theta_{0S}\| \leq 2(j+1)r_n\}$; we can choose this net in a way that $N_{S,j} \leq C^{|S|}$ for some constant C (see, for example, Lemma 5.2 of [30]). Letting $\theta^{S,j,i}_S = \phi^{S,j,i}$ and $\theta^{S,j,i}_S = 0$ for $k \in S^c$, we show this collection indeed forms a $2jr_n$ net of $\Theta_{S,j,n}$. To that end, fix $\theta \in \Theta_{S,j,n}$. Clearly, $\|\theta_S - \theta_{0S}\| \leq 2(j+1)r_n$. Find $1 \leq i \leq N_{S,j}$ such that $\|\theta^{S,j,i}_S - \theta_S\| \leq jr_n$. Then,

$$\|\theta^{S,j,i} - \theta\|_2^2 = \|\theta^{S,j,i}_S - \theta_S\|_2^2 + \|\theta_{S^c}\|_2^2 \leq (jr_n)^2 + (n - q_n)r_n^2/n^2 \leq 4j^2r_n^2,$$

proving our claim. Therefore, the union of balls $B_{S,j,i}$ of radius $2jr_n$ centered at $\theta^{S,j,i}$ for $1 \leq i \leq N_{S,j}$ cover $\Theta_{S,j,n}$. Since $E_{\theta_0} \mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n \mid y^{(n)}) \rightarrow 0$ by Theorem 3.4, it is enough to work with $E_{\theta_0} \mathbb{P}(\|\theta - \theta_0\|_2 > 2Mr_n, \text{supp}_{\delta_n}(\theta) \in \mathcal{S}_n \mid y^{(n)})$. Using the standard testing argument for establishing posterior convergence rates (see, for example, the proof of Proposition 5.1 in [9]), we arrive at

$$E_{\theta_0} \mathbb{P}(\|\theta - \theta_0\|_2 > 2Mr_n, \text{supp}_{\delta_n}(\theta) \in \mathcal{S}_n | y^{(n)}) I_{\mathcal{A}_n} \leq \sum_{S \in \mathcal{S}_n} \sum_{j \geq M} \sum_{i=1}^{N_{S,j}} 2 \sqrt{\beta_{j,S,i}} e^{-Cj^2 r_n^2}, \quad (16)$$

where

$$\beta_{S,j,i} = \frac{\mathbb{P}(B_{S,j,i})}{e^{-4r_n^2} \mathbb{P}(\|\theta - \theta_0\|_2 < 2r_n)}.$$

Let $r_n^2 = q_n \log n$. The proof of Theorem 3.1 is completed by deriving an upper bound to $\beta_{S,j,i}$ in the following Lemma 6.1 akin to Lemma 5.4 in [9].

Lemma 6.1. $\log \beta_{S,j,i} \leq |S| \log(2j) + C(|S| + |S_0|) \log n + C' r_n^2$.

Proof.

$$\beta_{S,j,i} \leq \frac{\mathbb{P}(\theta \in \mathbb{R}^n : \text{supp}_{\delta_n}(\theta) = S, \|\theta_S - \tilde{\theta}_S^{S,j,i}\|_2 < 2jr_n)}{e^{-4r_n^2} \mathbb{P}(\theta \in \mathbb{R}^n : \|\theta - \theta_0\|_2 < 2r_n)} \leq \frac{\mathbb{P}(\theta \in \mathbb{R}^n : |\theta_j| \leq \delta_n \forall j \in S^c, |\theta_j| > \delta_n \forall j \in S, \|\theta_S - \tilde{\theta}_S^{S,j,i}\|_2 < 2jr_n)}{e^{-4r_n^2} \mathbb{P}(\theta \in \mathbb{R}^n : \|\theta_{S_0} - \theta_{0S_0}\|_2 < r_n, \|\theta_{0S_0^c}\|_2 < r_n)} \leq \frac{e^{4r_n^2} \mathbb{P}}{e^{-4r_n^2} \mathbb{P}}$$

Next, we find an upper bound to

$$R_{S,j,i} = \frac{\mathbb{P}(|\theta_j| > \delta_n \forall j \in S, \|\theta_S - \tilde{\theta}_S^{S,j,i}\|_2 < 2jr_n)}{\mathbb{P}(\|\theta_{S_0} - \theta_{0S_0}\|_2 < r_n)}. \quad (18)$$

Let $v_q(r)$ denote the q -dimensional Euclidean ball of radius r centered at zero and $|v_q(r)|$ denote its volume. For the sake of brevity, denote $v_q = |v_q(1)|$, so that $|v_q(r)| = r^q v_q$. The numerator of (18) can be clearly bounded above by $|v_{|S|}(2jr_n)| \sup_{|\theta_j| > \delta_n \forall j \in S} \Pi_S(\theta_S)$. Since the set $\{\|\theta_{S_0} - \theta_{0S_0}\|_2 < r_n\}$ is contained in the ball $v_{|S_0|}(\|\theta_{0S_0}\|_2 + r_n) = \{\|\theta_{S_0}\|_2 \leq \|\theta_{0S_0}\|_2 + r_n\}$ and $\|\theta_{0S_0}\|_2 = \|\theta_0\|_2$, the denominator of (18) can be bounded below by $|v_{|S_0|}(r_n)| \inf_{v_{|S_0|}(t_n)} \Pi_{S_0}(\theta_{S_0})$, where $t_n = \|\theta_0\|_2 + r_n$. Putting together these inequalities and invoking Lemma 3.2, we have

$$R_{S,j,i} \leq \frac{(2jr_n)^{|S|} v_{|S|} \exp\{C|S| \log(1/\delta_n)\}}{r_n^{|S_0|} v_{|S_0|} \exp\{-C(|S_0| \log n + |S_0|^{3/4} t_n^{1/2})\}}. \quad (19)$$

Using $v_q \asymp (2\pi e)^{q/2} q^{-q/2-1/2}$ (see Lemma 5.3 in [9]) and $r_n^2 \geq q_n = |S_0|$, we can bound

$\log\{r_n^{|S|} v_{|S|} / (r_n^{|S_0|} v_{|S_0|})\}$ from above by $C(|S| \log n + r_n^2)$. Therefore, we have

$$\log R_{S,j,i} \leq |S| \log(2j) + C\{|S| \log n + r_n^2 + |S_0| \log n + |S| \log(1/\delta_n) + |S_0|^{3/4} t_n^{1/2}\}. \quad (20)$$

Now, since $\|\theta_0\|_2^2 \leq q_n \log^4 n$ and $r_n^2 = q_n \log n$, we have $t_n \lesssim q_n^{1/2} \log^2 n$ and hence

$|S_0|^{3/4} t_n^{1/2} \lesssim q_n \log n = r_n^2$. Substituting in (20), we have

$$\log R_{S,j,i} \leq |S| \log(2j) + C(|S| + |S_0|) \log n + C' r_n^2.$$

Finally, $\mathbb{P}(|\theta_1| < \delta_n)^{n-|S|} / \mathbb{P}(|\theta_1| < \delta_n)^{n-q_n} = \mathbb{P}(|\theta_1| < \delta_n)^{-|S|}$. Using Lemma 3.3, $\mathbb{P}(|\theta_1| < \delta_n) = (1 - \log n/n)$, which implies that $\mathbb{P}(|\theta_1| < \delta_n)^{-|S|} = e^{\log n}$.

Substituting the upper bound for $\beta_{S,j,i}$ obtained in Lemma 6.1, and noting that $|S| \leq Aq_n$ and $|N_{S,j}| \leq e^{C|S|}$, the expression in the left hand side of (16) can be bounded above by

$$\begin{aligned} & 2 \sum_{S \in \mathcal{S}_n} \sum_{j \geq M} \sum_{i=1}^{N_{S,j}} \exp\{Aq_n \log(2j)/2 \\ & \quad + C/2(A \\ & \quad + 1)q_n \log n + C' 2r_n^2/2\} e^{-Cj^2 r_n^2} \leq 2 \sum_{S \in \mathcal{S}_n} \sum_{j \geq M} \exp\{Cq_n + Aq_n \log(2j)/2 \\ & \quad + C/2(A \\ & \quad + 1)q_n \log n \\ & \quad + C' r_n^2/2\} e^{-Cj^2 r_n^2}. \end{aligned}$$

Since $|\mathcal{S}_n| \leq Aq_n \binom{n}{Aq_n} \leq Aq_n e^{Aq_n \log(ne/Aq_n)}$, it follows that $E_{\theta_0} \mathbb{P}(\theta : \|\theta - \theta_0\|_2 > 2Mr_n, \text{supp}_{\delta_n}(\theta) \in \mathcal{S}_n | y^{(n)}) \rightarrow 0$ for large $M > 0$.

When $a_n = n^{-(1+\beta)}$, the conclusion of Lemma 6.1 remains unchanged and the proof of Theorem 3.4 does not require $q_n \gtrsim \log n$. The rest of the proof remains exactly the same.

ACKNOWLEDGEMENTS

We thank Dr. Ismael Castillo and Dr. James Scott for sharing source code. Dr. Anirban Bhattacharya, Dr. Debdeep Pati and Dr. Natesh S. Pillai acknowledge support from the Office of Naval Research (ONR BAA 14-0001). Dr. Pillai is also partially supported by NSF-DMS 1107070. Dr. David B. Dunson is partially supported by the DARPA MSEE program and the grant number R01 ES017240-01 from the National Institute of Environmental Health Sciences (NIEHS) of the National Institutes of Health (NIH).

APPENDIX

Proof of Proposition 2.1

When $a = 1/n$, $\phi_j \sim \text{Beta}(1/n, 1 - 1/n)$ marginally. Hence, the marginal distribution of θ_j given τ is proportional to

$$\int_{\phi_j=0}^1 e^{-|\theta_j|/(\phi_j \tau)} \left(\frac{\phi_j}{1 - \phi_j} \right)^{1/n} \phi_j^{-2} d\phi_j.$$

Substituting $z = \phi_j/(1 - \phi_j)$ so that $\phi_j = z/(1 + z)$, the above integral reduces to

$$e^{-|\theta_j|/\tau} \int_{z=0}^{\infty} e^{-|\theta_j|/(\tau z)} z^{-(2-1/n)} dz \propto e^{-|\theta_j|/\tau} |\theta_j|^{1/n-1}.$$

In the general case, $\varphi_j \sim \text{Beta}(a, (n-1)a)$ marginally. Substituting $z = \varphi_j/(1 - \varphi_j)$ as before, the marginal density of θ_j is proportional to

$$e^{-|\theta_j|/\tau} \int_{z=0}^{\infty} e^{-|\theta_j|/(\tau z)} z^{-(2-a)} \left(\frac{1}{1+z} \right)^{na-1} dz.$$

The above integral can clearly be bounded below by a constant multiple of

$$e^{-|\theta_j|/\tau} \int_{z=0}^1 e^{-|\theta_j|/(\tau z)} z^{-(2-a)} dz.$$

The above expression clearly diverges to infinity as $|\theta_j| \rightarrow 0$ by the monotone convergence theorem.

Proof of Theorem 2.2

Integrating out τ , the joint posterior of $\varphi \mid \theta$ has the form

$$\pi(\phi_1, \dots, \phi_{n-1} \mid \theta) \propto \prod_{j=1}^n \left[\phi_j^{a-1} \frac{1}{\phi_j} \right] \int_{\tau=0}^{\infty} e^{-\tau/2} \tau^{\lambda-n-1} e^{-\sum_{j=1}^n |\theta_j|/(\phi_j \tau)} d\tau. \quad (\text{A.1})$$

We now state a result from the theory of normalized random measures (see, for example, (36) in [22]). Suppose $T_1, \dots, T_n$ are independent random variables with T_j having a density f_j on $(0, \infty)$. Let $\varphi_j = T_j/T$ with $T = \sum_{j=1}^n T_j$. Then, the joint density f of $(\varphi_1, \dots, \varphi_{n-1})$ supported on the simplex $\mathcal{S}^{n-1}$ has the form

$$f(\phi_1, \dots, \phi_{n-1}) = \int_{t=0}^{\infty} t^{n-1} \prod_{j=1}^n f_j(\phi_j t) dt, \quad (\text{A.2})$$

where $\phi_n = 1 - \sum_{j=1}^{n-1} \phi_j$. Setting $f_j(x) \propto \frac{1}{x^\delta} e^{-|\theta_j|/x} e^{-x/2}$ in (A.2), we get

$$f(\phi_1, \dots, \phi_{n-1}) = \left[\prod_{j=1}^n \frac{1}{\phi_j^\delta} \right] \int_{t=0}^{\infty} e^{-t/2} t^{n-1-n\delta} e^{-\sum_{j=1}^n |\theta_j|/(\phi_j t)} dt. \quad (\text{A.3})$$

We aim to equate the expression in (A.3) with the expression in (A.1). Comparing the exponent of φ_j gives us $\delta = 2 - a$. The other requirement $n - 1 - n\delta = \lambda - n - 1$ is also satisfied, since $\lambda = na$. The proof is completed by observing that f_j corresponds to a $\text{giG}(a - 1, 1, 2|\theta_j|)$ when $\delta = 2 - a$.

Proof of Proposition 3.1

By (10),

$$\Pi(\theta_j) = \frac{(1/2)^a}{2\Gamma(a)} \int_{\psi_j=0}^{\infty} e^{-|\theta_j|/\psi_j} \psi_j^{a-2} e^{-\psi_j/2} d\psi_j = \frac{(1/2)^a}{2\Gamma(a)} \int_{z=0}^{\infty} e^{-z|\theta_j|} z^{-a} e^{-2/z} dz.$$

The result follows from 8.432.7 in [15].

Proof of Lemma 3.2

Letting $h(x) = \log \Pi(x)$, we have $\log \Pi_S(\eta) = \sum_{j=1}^{|S|} h(\eta_j)$.

We first prove (13). Since $\Pi(x)$, and hence $h(x)$, is monotonically decreasing in $|x|$, and $|\eta_j| > \delta$ for all j , we have $\log \Pi_S(\eta) \leq |S|h(\delta)$. Using $K_a(z) \asymp z^{-a}$ for $|z|$ small and $\Gamma(a) \asymp a^{-1}$ for a small, we have from (11) that $\Pi(\delta) \asymp a^{-1}|\delta|^{(a-1)}$ and hence $h(\delta) \asymp (1-a) \log(\delta^{-1}) - \log a^{-1} + C = C \log(\delta^{-1})$.

We next prove (14). Noting that $K_a(z) \gtrsim e^{-z}/z$ for $|z|$ large (section 9.7 of [1]), we have from (11) that $-h(x) \leq \log a^{-1} + 3/2 \log|x| + \sqrt{2} \sqrt{|x|}$ for $|x|$ large. Using Cauchy-Schwartz

inequality twice, we have $(\sum_{j=1}^{|S|} \sqrt{|\eta_j|})^4 \leq |S|^3 \|\eta\|_2^2$, which implies $\sum_{j=1}^{|S|} \sqrt{|\eta_j|} \leq |S|^{3/4} \|\eta\|_2^{1/2} \leq |S|^{3/4} m^{1/2}$.

Proof of Lemma 3.3

Using the representation (9), we have $\mathbb{P}(|\theta_1| > \delta \mid \psi_1) = e^{-\delta/\psi_1}$, so that,

$$\begin{aligned} \mathbb{P}(|\theta_1| > \delta) &= \frac{(1/2)^a}{\Gamma(a)} \int_0^{\infty} e^{-\delta/x} x^{a-1} e^{-x/2} dx \\ &= \frac{(1/2)^a}{\Gamma(a)} \left\{ \int_0^{4\delta} e^{-\delta/x} x^{a-1} e^{-x/2} dx \right. \\ &\quad \left. + \int_{4\delta}^{\infty} e^{-\delta/x} x^{a-1} e^{-x/2} dx \right\} \leq \frac{(1/2)^a}{\Gamma(a)} \left\{ C + \int_{4\delta}^{\infty} \frac{e^{-x/2}}{x} dx \right\} \leq \frac{(1/2)^a}{\Gamma(a)} \left\{ C \right. \\ &\quad \left. + \int_{2\delta}^{\infty} \frac{e^{-t}}{t} dt \right\}, \end{aligned} \tag{A.4}$$

where $C > 0$ is a constant independent of δ . Using a bound for the incomplete gamma function from Theorem 2 of [2],

$$\int_{2\delta}^{\infty} \frac{e^{-t}}{t} dt \leq -\log(1 - e^{-2\delta}) \leq -\log(\delta), \tag{A.5}$$

for δ small. The proof is completed by noting that $(1/2)^a$ is bounded above by a constant and $C + \log(1/\delta) \leq 2 \log(1/\delta)$ for δ small enough.

Proof of Theorem 3.4

For $\theta \in \mathbb{R}^n$, let $f_\theta(\cdot)$ denote the probability density function of a $N_n(\theta, I_n)$ distribution and f_{θ_i} denote the univariate marginal $N(\theta_i, 1)$ distribution. Let $S_0 = \text{supp}(\theta_0)$. Since $|S_0| = q_n$, it suffices to prove that

$$\lim_{n \rightarrow \infty} E_{\theta_0} \mathbb{P}(|\text{supp}_{\delta_n}(\theta) \cap S_0^c| > Aq_n | y^{(n)}) \rightarrow 0.$$

Let $\mathcal{B}_n = \{|\text{supp}_{\delta_n}(\theta) \cap S_0^c| > Aq_n\}$. By (10), $\{\theta_i, i \in S_0^c\}$ is independent of $\{\theta_i, i \in S_0\}$ conditionally on $y^{(n)}$. Hence

$$\mathbb{P}(\mathcal{B}_n | y^{(n)}) = \frac{\int_{\mathcal{B}_n} \prod_{i \in S_0^c} \frac{f_{\theta_i}(y_i)}{f_0(y_i)} d\Pi(\theta_i)}{\int \prod_{i \in S_0^c} \frac{f_{\theta_i}(y_i)}{f_0(y_i)} d\Pi(\theta_i)} := \frac{\mathcal{N}'_n}{\mathcal{D}'_n}, \quad (\text{A.6})$$

where $\mathcal{N}'_n$ and $\mathcal{D}'_n$ respectively denote the numerator and denominator of the expression in (A.6). Observe that

$$E_{\theta_0} \mathbb{P}(\mathcal{B}_n | y^{(n)}) \leq E_{\theta_0} \mathbb{P}(\mathcal{B}_n | y^{(n)}) 1_{\mathcal{A}'_n} + P_{\theta_0}(\mathcal{A}'_n{}^c), \quad (\text{A.7})$$

where $\mathcal{A}'_n$ is a subset of $\sigma(y^{(n)})$ as in Lemma 5.2 of [9] (replacing θ by $\theta_{S_0^c}$ and θ_0 by 0) defined as

$$\mathcal{A}'_n = \{\mathcal{D}'_n \geq e^{-r_n^2} \mathbb{P}(\|\theta_{S_0^c}\|_2 \leq r_n)\},$$

with $P_{\theta_0}(\mathcal{A}'_n{}^c) \leq e^{-r_n^2}$ for some sequence of positive real numbers r_n . We set $r_n^2 = q_n$ here. With this choice, from (A.7),

$$E_{\theta_0} \mathbb{P}(\mathcal{B}_n | y^{(n)}) \leq \frac{\mathbb{P}(\mathcal{B}_n)}{e^{-r_n^2} \mathbb{P}(\|\theta_{S_0^c}\|_2 \leq r_n)} + e^{-r_n^2}. \quad (\text{A.8})$$

We have $\mathbb{P}(\|\theta_{S_0^c}\|_2 \leq r_n) \geq \mathbb{P}(|\theta_j| < r_n / \sqrt{n} \forall j \in S_0^c) = \mathbb{P}(|\theta_1| < r_n / \sqrt{n})^{n-q_n}$, with the equality following from the representation in (10). Using Lemma 3.3,

$\mathbb{P}(|\theta_1| < r_n / \sqrt{n}) \geq 1 - \log n / n$, implying $\mathbb{P}(\|\theta_{S_0^c}\|_2 \leq r_n) \geq e^{-C \log n}$. Next, clearly $\mathbb{P}(\mathcal{B}_n) = \mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n)$. As indicated in Section 3.1, $|\text{supp}_{\delta_n}(\theta)| \sim \text{Binomial}(n, \zeta_n)$, with $\zeta_n = \mathbb{P}(|\theta_1| > \delta_n) \sim \log n / n$ in view of Lemma 3.3. A version of Chernoff's inequality for the binomial distribution [17] states that for $B \sim \text{Binomial}(n, \zeta)$ and $\zeta < a < 1$,

$$\mathbb{P}(B > an) \leq \left\{ \left(\frac{\zeta}{a} \right)^a e^{a-\zeta} \right\}^n. \quad (\text{A.9})$$

When $a_n = 1/n$, $q_n \sim C_0 \log n$ for some constant $C_0 > 0$. Setting $a_n = Aq_n/n$, clearly $\zeta_n < a_n$ for some $A > 1/C_0$. Substituting in (A.9), we have $\mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n) \leq e^{Aq_n \log(e \log n) - Aq_n \log Aq_n}$. Choosing $A = 2e/C_0$ and using the fact that $q_n \sim C_0 \log n$, we

obtain $\mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n) \leq e^{-Aq_n \log 2}$. Substituting the bounds for $\mathbb{P}(\mathcal{B}_n)$ and $\mathbb{P}(\|\theta_{s_0^c}\|_2 \leq r_n)$ in (A.8) and choosing larger A if necessary, the expression in (A.7) goes to zero.

If $a_n = n^{-(1+\beta)}$, $\zeta_n \sim \log n/n^{1+\beta}$ in view of Lemma 3.3. In (A.9), set $a_n = Aq_n/n$ as before. Clearly $\zeta_n < a_n$. Substituting in (A.9), we have $\mathbb{P}(|\text{supp}_{\delta_n}(\theta)| > Aq_n) \leq e^{-CAq_n \log n}$.

Substituting the bounds for $\mathbb{P}(\mathcal{B}_n)$ and $\mathbb{P}(\|\theta_{s_0^c}\|_2 \leq r_n)$ in (A.8), the expression in (A.7) goes to zero.

REFERENCES

1. Abramowitz, M.; Stegun, IA. Handbook of mathematical functions: with formulas, graphs, and mathematical tables. Vol. 55. Dover publications; 1965.
2. Alzer H. On some inequalities for the incomplete gamma function. *Mathematics of Computation*. 1997; 66(218):771–778.
3. Armagan A, Dunson D, Lee J. Generalized double pareto shrinkage. *Statistica Sinica*, to appear. 2013
4. Armagan, Artin; Dunson, David B.; Lee, Jaeyong; Bajwa, Waheed U.; Strawn, Nate. Posterior consistency in linear models under shrinkage priors. *Biometrika*. 2013; 100(4):1011–1018.
5. Benjamini, Yoav; Hochberg, Yosef. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1995:289–300.
6. Bontemps D. Bernstein–von mises theorems for gaussian regression with increasing number of regressors. *The Annals of Statistics*. 2011; 39(5):2557–2584.
7. Carvalho CM, Polson NG, Scott JG. Handling sparsity via the horseshoe. *Journal of Machine Learning Research W&CP*. 2009; 5(73–80)
8. Carvalho CM, Polson NG, Scott JG. The horseshoe estimator for sparse signals. *Biometrika*. 2010; 97(2):465–480.
9. Castillo I, van der Vaart A. Needles and straws in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics*. 2012; 40(4):2069–2101.
10. Castillo, Ismael; Schmidt-Hieber, Johannes; van der Vaart, Aad W. Bayesian linear regression with sparse priors. *arXiv preprint arXiv: 1403.0735*. 2014
11. Donoho DL, Johnstone IM, Hoch JC, Stern AS. Maximum entropy and the nearly black object. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1992:41–81.
12. Efron, Bradley. Microarrays, empirical bayes and the two-groups model. *Statistical Science*. 2008:1–22.
13. Efron, Bradley. Large-scale inference: empirical Bayes methods for estimation, testing, and prediction. Vol. 1. Cambridge University Press; 2010.
14. Ghosal S. Asymptotic normality of posterior distributions in high-dimensional linear models. *Bernoulli*. 1999; 5(2):315–331.
15. Gradshteyn IS, Ryzhik IM. Corrected and enlarged edition. *Tables of Integrals, Series and Products* Academic Press, New York. 1980
16. Griffin JE, Brown PJ. Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis*. 2010; 5(1):171–188.
17. Hagerup, Torben; Rüb, Christine. A guided tour of chernoff bounds. *Information processing letters*. 1990; 33(6):305–308.
18. Hans C. Elastic net regression modeling with the orthant normal prior. *Journal of the American Statistical Association*. 2011; 106(496):1383–1393.
19. Johnson, Valen E.; Rossell, David. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*. 2012; 107(498):649–660.

20. Johnstone IM, Silverman BW. Needles and straw in haystacks: Empirical bayes estimates of possibly sparse sequences. *The Annals of Statistics*. 2004; 32(4):1594–1649.
21. Johnstone IM, Silverman BW. EBayesthresh: R and s-plus programs for empirical bayes thresholding. *J. Statist. Soft.* 2005; 12:1–38.
22. Kruijer W, Rousseau J, van der Vaart A. Adaptive bayesian density estimation with location-scale mixtures. *Electronic Journal of Statistics*. 2010; 4:1225–1257.
23. Leng C. Variable selection and coefficient estimation via regularized rank regression. *Statistica Sinica*. 2010; 20(1):167.
24. Naidu Narisetty, Naveen; He, Xuming, et al. Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics*. 2014; 42(2):789–817.
25. Park T, Casella G. The bayesian lasso. *Journal of the American Statistical Association*. 2008; 103(482):681–686.
26. Polson, NG.; Scott, JG. Shrink globally, act locally: Sparse Bayesian regularization and prediction. In: Bernardo, JM.; Bayarri, MJ.; Berger, JO.; Dawid, AP.; Heckerman, D.; Smith, AFM.; West, M., editors. *Bayesian Statistics 9*. New York: Oxford University Press; 2010. p. 501-538.
27. Scott JG, Berger JO. Bayes and empirical-bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*. 2010; 38(5):2587–2619.
28. Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1996:267–288.
29. Van De Geer, Sara A.; Bühlmann, Peter, et al. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*. 2009; 3:1360–1392.
30. Vershynin R. Introduction to the non-asymptotic analysis of random matrices. Arxiv preprint arxiv: 1011.3027. 2010
31. Wang H, Leng C. Unified lasso estimation by least squares approximation. *Journal of the American Statistical Association*. 2007; 102(479):1039–1048.
32. West M. On scale mixtures of normal distributions. *Biometrika*. 1987; 74(3):646–648.
33. Zhou M, Carin L. Negative binomial process count and mixture modeling. arXiv preprint arXiv: 1209.3442. 2012
34. Zou H. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*. 2006; 101(476):1418–1429.

Figure 1.
Marginal density of the DL prior with $a = 1/2$ in comparison to other shrinkage priors.

Figure 2.

Simulation results from a single replicate with $n = 200$, $q_n = 10$, $A = 7$. Blue circles = entries of y , red circles = posterior median of θ , shaded region: 95% point wise credible interval for θ . Left panel: Bayesian lasso, right panel: $DL_{1/n}$ prior

Figure 3.

Simulation results from a single replicate with $n = 200$, $q_n = 10$, $A = 7$. Blue circles = entries of y , red circles = posterior median of θ , shaded region: 95% point wise credible interval for θ . Left panel: Horseshoe, right panel: $DL_{1/2}$ prior

Figure 4.
Histogram of z-values

Table 1

Squared error comparison over 100 replicates. Average squared error across replicates reported for DL (Dirichlet–Laplace) with $a = 1/n$ and $a = 1/2$, LS (Lasso), EBMed (Empirical Bayes median), PM (Point mass prior), BL (Bayesian lasso) and HS (horseshoe). Signal strength $A = 5, 6$.

$\frac{Q_{in} \sigma^2}{\Pi}$	100						200					
	5		10		20		5		10		20	
A	5	6	5	6	5	6	5	6	5	6	5	6
DL _{1/n}	21.90	14.09	32.43	21.81	55.26	38.71	44.86	31.57	50.24	35.39	95.14	85.54
DL _{1/2}	14.77	12.54	21.79	18.26	37.13	31.06	28.52	25.46	43.74	37.45	75.37	65.29
LS	22.21	20.67	36.41	36.54	65.91	66.60	43.37	43.51	74.05	75.98	139.20	137.94
EBMed	15.42	13.68	28.26	27.86	57.93	58.00	28.94	27.50	56.03	57.65	120.67	119.85
PM	14.50	11.99	24.97	23.66	49.92	49.16	26.66	24.86	49.96	50.89	103.40	101.69
BL	30.84	31.61	44.32	46.65	61.23	63.46	59.70	62.30	88.10	94.02	125.89	128.72
HS	12.50	9.31	21.63	17.63	39.60	35.89	23.41	18.75	42.95	38.21	81.07	74.31

Table 2

Squared error comparison over 100 replicates. Average squared error across replicates reported for DL (Dirichlet–Laplace) with $a = 1/n$ and $a = 1/2$, LS (Lasso), EBMed (Empirical Bayes median), PM (Point mass prior), BL (Bayesian lasso) and HS (horseshoe). Signal strength $A = 7, 8$.

$\frac{Q_{in} \sigma^2}{\Pi}$	100						200					
	5		10		20		5		10		20	
A	7	8	7	8	7	8	7	8	7	8	7	8
DL _{1/n}	8.20	7.19	17.29	15.35	32.00	29.40	16.07	14.28	33.00	30.80	65.53	59.61
DL _{1/2}	11.77	11.87	17.57	18.36	30.24	30.42	20.30	22.70	36.51	35.30	61.43	65.79
LS	21.25	19.09	38.68	37.25	68.97	69.05	41.82	41.18	75.55	75.12	137.21	136.25
EBMed	13.64	12.47	29.73	27.96	60.52	60.22	26.10	25.52	57.19	56.05	119.41	119.35
PM	12.15	10.98	25.99	24.59	51.36	50.98	22.99	22.26	49.42	48.42	101.54	101.62
BL	33.05	33.63	49.85	50.04	68.35	68.54	64.78	69.34	99.50	103.15	133.17	136.83
HS	8.30	7.93	18.39	16.27	37.25	35.18	15.80	15.09	35.61	33.58	72.15	70.23

Table 3

Squared error comparison over 100 replicates. Average squared error for the posterior median reported for BL (Bayesian Lasso), HS (horseshoe) and DL (Dirichlet–Laplace) with $a = 1/n$ and $a = 1/2$ respectively.

		1000					
n	A	2	3	4	5	6	7
BL		299.30	385.68	424.09	450.20	474.28	493.03
HS		306.94	353.79	270.90	205.43	182.99	168.83
DL _{1/n}		368.45	679.17	671.34	374.01	213.66	160.14
DL _{1/2}		267.83	315.70	266.80	213.23	192.98	177.20