

Continuous shrinkage prior revisited: a collapsing behavior and remedy

Se Yoon Lee, Debdeep Pati, and Bani K. Mallick

Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX, U.S.A.

July 7, 2020

Abstract

Modern genomic studies are increasingly focused on identifying more and more genes clinically associated with a health response. Commonly used Bayesian shrinkage priors are designed primarily to detect only a handful of signals when the dimension of the predictors is very high. In this article, we investigate the performance of a popular continuous shrinkage prior in the presence of relatively large number of true signals. We draw attention to an undesirable phenomenon; the posterior mean is rendered very close to a null vector, caused by a sharp underestimation of the global-scale parameter. The phenomenon is triggered by the absence of a tail-index controlling mechanism in the Bayesian shrinkage priors. We provide a remedy by developing a global-local-tail shrinkage prior which can automatically learn the tail-index and can provide accurate inference even in the presence of moderately large number of signals. The collapsing behavior of the Horseshoe with its remedy is exemplified in numerical examples and in two gene expression datasets.

KEYWORDS: The Horseshoe; Gene expression data; Collapsing behavior; The GLT prior.

1. INTRODUCTION

Development of sophisticated data acquisition techniques in gene expression microarray, among many other fields triggered the development of innovative statistical methods [7, 18, 22] to identify relevant predictors associated with a response out of a large number of predictors, but only with a smaller number of samples. This large p , small n paradigm is arguably the most researched topic in the last decade.

Focusing on the linear regression model for simplicity, consider responses y_i corresponding to p dimensional covariates x_i . Let $y = (y_1, \dots, y_n)^T$ and X denote the $n \times p$ covariate matrix with x_i as the rows and

$$y = X\beta + \sigma\epsilon, \quad \epsilon \sim \mathcal{N}_n(0, I_n).$$

When $p \gg n$, a natural assumption for a meaningful inference on β is to assume sparsity, i.e., substantially many coefficients in β are assumed to be zeros or approximately zeros. This was believed to be a reasonable assumption in the context of gene-expression studies [16] where it is believed that only a fraction of genes is really responsible for affecting the response. The true non-zero coefficients are referred to as the signal coefficients and the remaining are noise coefficients. Throughout the remainder of the paper, we make use of the following definition of the sparsity level. Let q denote the number of true non-zero signals in β : then the following ratio $s = q/p$ is referred to as the sparsity level

$$s = \frac{q}{p} = \frac{\text{the number of relevant predictors}}{\text{total number of predictors}}. \quad (1)$$

Statisticians have devised a number of penalized regression techniques for estimating β in $p \gg n$ setting under the assumption of sparsity [22]. From a Bayesian point view, sparsity favoring mixture priors with separate control on the signal and noise coefficients have been proposed [19, 26, 31, 51]. Although they often lead to attractive theoretical properties [10, 11], computational issues and considerations that many of the β_j 's may be small but not exactly zero has led to a rich

variety of continuous shrinkage priors being proposed recently [8, 9, 21, 33, 43], which can be unified through a global-local (GL) scale mixture representation of [37]. Among the continuous shrinkage priors, the Horseshoe [8, 9] is possibly the most visible and acclaimed method.

In this paper, we revisit the performance of the Horseshoe prior on real applications involving gene-expression studies. It is important to point out that the posterior obtained using Horseshoe has remarkable finite sample performance and enjoys several optimal theoretical properties [46, 47, 48, 49] when the underlying sparsity level s in (1) is very small in the high dimensional setting $p \geq n$. However, in many applications as outlined in Section 3, it is necessary to consider a moderately sparse regime with a relatively higher value of s to reflect the fact that there are many small signals, possibly with a strong correlation amongst corresponding covariates. This is especially true in cancer studies where the number of interesting genes, i.e., q in (1) is growing in commensurate with the denominator p in (1) which itself has grown enormously due to the success of the Human Genome Project. For example, BRCA1 and BRCA2 are linked to the risk of breast cancer risk over 20 years ago. However, ‘beyond BRCA1 & 2 movement’ [38] has led to discovery of more interesting protein coding genes [41, 42] with the latest estimate of the protein-coding gene count being $p = 21,306$ [50].

In this paper, we empirically showed that when the true sparsity level is moderate, the Horseshoe may end up estimating β by an approximately null vector. We refer to this as a collapsing behavior typically caused by underestimation of the global-scale parameters; a similar phenomenon is also observed by [2]. We demonstrate this in a wide variety of simulation examples and in popular gene expression datasets. From a theoretical standpoint, we provide justification to this collapsing behavior by examining the tail behavior of the Horseshoe. In particular, the fixed tail-index limits the flexibility to estimate a moderately large number of signals resulting in very small estimates of the global-scale parameter. We propose a remedy for this collapsing behavior by introducing a tail-controlling mechanism within the Horseshoe formulation, called the global-local-tail shrinkage priors. We demonstrate superior performance of the prior in simulated and real examples.

2. A REVIEW OF CONTINUOUS SHRINKAGE PRIORS

A popular computationally scalable class of continuous shrinkage priors can be represented as global-local scale mixtures of Gaussian distribution:

$$\beta_j | \lambda_j, \tau, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2 \tau^2 \sigma^2), \quad \sigma^2 \sim h(\sigma^2), \quad (j = 1, \dots, p), \quad (2)$$

$$\lambda_j \sim f(\lambda_j), \quad \tau \sim g(\tau), \quad (j = 1, \dots, p), \quad (3)$$

where f , g , and h are densities supported on $(0, \infty)$. The top-level scale parameters, $\{\lambda_j\}_{j=1}^p$ and τ in (2) – (3), are referred to as the local-scale parameters and the global-scale parameter, respectively. Different choices of f and g for the top-level scale parameters lead to different class of priors [4]. Ideally, g should have a substantial mass near zero to enforce shrinkage towards zero which is compensated by allowing f to have heavier tails in order to capture the large signals and prevent over-shrinkage. In the high-dimensional setting, the choices of f and g play a key role in controlling the effective sparsity and concentration of the prior and posterior distributions [1, 29, 34, 37, 40, 53]. Choosing a half-Cauchy distribution, $\pi(x) = \mathcal{C}^+(x|0, 1) \propto 1/(1 + x^2)$, for f and g in (3) leads to the Horseshoe [9]. Under the sparsity assumption $s \rightarrow 0$ as $n, p \rightarrow \infty$, it is known that the posterior mean of Horseshoe, the Horseshoe estimator, possesses many nice theoretical properties [1, 4, 37, 40, 48]. For instance, the Horseshoe estimator is robust and attains the minimax-optimal rate for squared error loss up to a multiplicative constant under certain conditions [45, 49]. Highly scalable algorithms are recently proposed for the Horseshoe [5, 25].

3. COLLAPSING BEHAVIOR OF THE HORSESHOE ESTIMATOR IN GENE EXPRESSION DATA

In the following, we analyze the prostate cancer data [6, 14, 15, 16, 39], which has been widely adopted as a prototype real application in high dimensional sparse regression problems. A similar analysis with a breast cancer data is in §1.1 of the supplementary materials.

3.1 The Horseshoe applied to the prostate cancer data

The prostate cancer data is summarized in a matrix $X \in \mathfrak{R}^{102 \times 6033}$ comprising of two parts. The first 50 rows of X , $X[1 : 50, \cdot] \in \mathfrak{R}^{50 \times 6033}$, correspond to healthy controls, and the remaining rows, $X[51 : 102, \cdot] \in \mathfrak{R}^{52 \times 6033}$, correspond to cancer patients. The j -th column vector of X , $X[\cdot, j] \in \mathfrak{R}^{102}$, $j = 1, \dots, 6,033$, represents gene expression levels of the j -th gene.

The main goal of the study is to discover a small number of interesting genes whose expression levels differ between the two classes [16]. Such genes are then investigated for a causal link for the development of prostate cancer. First, a multiple testing procedure for this data is carried out as follows: For each $j = 1, \dots, 6,033$, a two-sample t -test statistic with 100 degrees of freedom is obtained based on $X[\cdot, j] \in \mathfrak{R}^{102}$, and the t -test statistics are converted to z -test statistics using quantile transformation $y_j = \Phi^{-1}(F_{100 \text{ d.f.}}(t_j))$, where $\Phi(\cdot)$ and $F_{100 \text{ d.f.}}(\cdot)$ are distribution functions of $\mathcal{N}_1(0, 1)$ and t_{100} , respectively; refer to Section 2.1 of [16]. The j -th null hypothesis H_{0j} posits no difference in the gene expression levels for the j -th gene between the healthy controls and cancer patients. If the global null hypothesis $\cap_{j=1}^{6033} H_{0j}$ is true, the histogram of $\{y_j\}_{j=1}^{6033}$ should mimic a standard normal density closely. The histogram of $\{y_j\}_{i=1}^{6033}$ along with the standard normal density is displayed in the left panel in the Figure 1. Presence of outliers, possibly corresponding to cancerous genes [15], is evident. As in [16], we convert the problem to an estimation problem where a p -dimensional vector $\beta \in \mathfrak{R}^p$, $p = 6,033$ is estimated from a sparse normal means model:

$$y = \beta + \sigma\epsilon, \quad \epsilon \sim \mathcal{N}_p(0, I_p) \quad (4)$$

where σ is unknown. We use the Horseshoe prior $\pi_{\text{HS}}(\beta)$ for β , and the Jeffrey's prior for σ^2 , $\pi(\sigma^2) \propto 1/\sigma^2$ [24]. For implementation, the R function `Horseshoe` within the R package `Horseshoe` is used. More specifically, we used `Horseshoe(y = y, X = X, method.tau = "halfCauchy", method.sigma = "Jeffreys", burn = 10000, nmc = 10000, thin = 100)` where $y = y$ and $X = I_p$, to produce 100 thinned realizations from the posterior distribution $\pi(\beta|y)$ via Markov chain Monte Carlo (MCMC).

To investigate the behavior of the Horseshoe as the number of genes used increases, we con-




Figure 1: Histogram of z -values $\{y_j\}_{j=1}^{6033}$ (left panel), and the scatter plot of $\{(y_j, \hat{\beta}_j)\}_{j=1}^p$ (right panel) obtained using Horseshoe when applied to dataset \mathcal{P}_l , $l = 1, 2, 3, 4$. Dotted line is $y = x$.

structured four datasets, $\mathcal{P}_1 = \{y_j\}_{j=1}^{p=50}$, $\mathcal{P}_2 = \{y_j\}_{j=1}^{p=100}$, $\mathcal{P}_3 = \{y_j\}_{j=1}^{p=200}$, and $\mathcal{P}_4 = \{y_j\}_{j=1}^{p=6033}$. Note that \mathcal{P}_4 is the full dataset, and \mathcal{P} stands for prostate. The results of posterior inference are shown on the right panel of the Figure 1. In the panel, we plotted p ordered pairs $\{(y_j, \hat{\beta}_j)\}_{j=1}^p$ such that $\hat{\beta}_j$ represents the posterior mean of β_j , which are obtained from each dataset \mathcal{P}_l , $l = 1, 2, 3$, and 4.

From the robustness property of Horseshoe, we expect to observe an ideal reverse- S -shape curve for $\{(y_j, \hat{\beta}_j)\}_{j=1}^p$ as in Figure 1 of [49]. However, the right panel of the Figure 1 shows that the robustness property is manifested only when $p = 50$, and the property disappears as p increases, and when $p = 200$ or more the posterior mean of $\beta \in \mathbb{R}^p$ is essentially rendered a p -dimensional zero vector.

4. UNDERSTANDING THE COLLAPSING BEHAVIOR OF THE HORSESHOE ESTIMATOR

We first elaborate on the generation of high-dimensional data $(y, X) \in \mathbb{R}^n \times \mathbb{R}^{n \times p}$ from a sparse linear regression (1) corresponding to a simulation environment associated with a setting $(n, p, q,$

ϱ , SNR),

$$(y, X) \sim p(y, X) = \mathcal{N}_n(y|X\beta_0, \sigma_0^2 I_n) \cdot \prod_{i=1}^n \mathcal{N}_p(X_i^T|0, \Upsilon(\varrho)), \quad \Upsilon(\varrho) = \varrho J_p + (1 - \varrho)I_p, \quad (5)$$

where $p(y, X)$ is a true data generating process, and β_0 is a p -dimensional true data generating parameter such that $\beta_{0,1} = \dots = \beta_{0,q} = 1$ and $\beta_{0,q+1} = \dots = \beta_{0,p} = 0$. The first q coefficients of β_0 are the true signals of unit signal strengths. I and J indicate an identity matrix and a matrix whose elements are ones, respectively. The signal-to-noise ratio (SNR) is defined by $\text{SNR} = \text{var}(X\beta_0)/\text{var}(\sigma_0\epsilon)$. The value ϱ is a number associated with column-wise correlations in the design matrix X .

We use the following three steps to generate data (y, X) according to (5). i) Generate a matrix $X \in \mathfrak{R}^{n \times p}$ where each row vector $X_i^T \in \mathfrak{R}^p$ is independently sampled from $\mathcal{N}_p(0, \Upsilon(\varrho))$. Next, center the matrix X column-wise so that each column vector $X[:, j] \in \mathfrak{R}^n$ ($j = 1, \dots, p$) has zero mean. Then, normalize each vector to make Euclidean norm to be one. ii) Generate n -dimensional Gaussian errors $\epsilon \sim \mathcal{N}_n(0, I_n)$. iii) Add the mean $X\beta_0$ and the error $\sigma_0\epsilon$ to create responses $y = X\beta_0 + \sigma_0\epsilon$, where $\sigma_0^2 = \text{var}(X\beta_0)/\{\text{SNR} \cdot \text{var}(\epsilon)\}$ and $\text{var}(z) = \sum_{i=1}^n (z_i - \bar{z})^2/(n - 1)$ for $z \in \mathfrak{R}^n$.

A single simulation environment $(n, p, q, \varrho, \text{SNR})$ corresponds to multiple replications of the datasets (y, X) . To investigate the behavior of the Horseshoe when the sparsity level s in (1) increases, we generated four artificial datasets $\mathcal{A}_l = (y, X) \in \mathfrak{R}^n \times \mathfrak{R}^{n \times p}$, $l = 1, 2, 3, 4$, corresponding to four simulation environments $(n = 100, p = 500, q, \varrho = 0, \text{SNR} = 5)$ such that $q = 2, 5, 8, 13$. Therefore, the sparsity levels of the four datasets \mathcal{A}_l ($l = 1, 2, 3, 4$) are $2/500 = 0.004$ (\mathcal{A}_1), $5/500 = 0.01$ (\mathcal{A}_2), $8/500 = 0.016$ (\mathcal{A}_3), and $13/500 = 0.026$ (\mathcal{A}_4), respectively.

The results of posterior inference are displayed in Figure 2. Panels are arranged in a way that the sparsity level increases from the left to right. Panels in the first, second, and third rows in the Figure 2 display the 95% credible intervals for $\{\beta_j\}_{j=1}^p$, and those of $\{\lambda_j\}_{j=1}^p$, and posterior correlations $\{\text{cor}(\lambda_j, \tau|y)\}_{j=1}^p$, respectively. For the ease of visualization, results corresponding to only

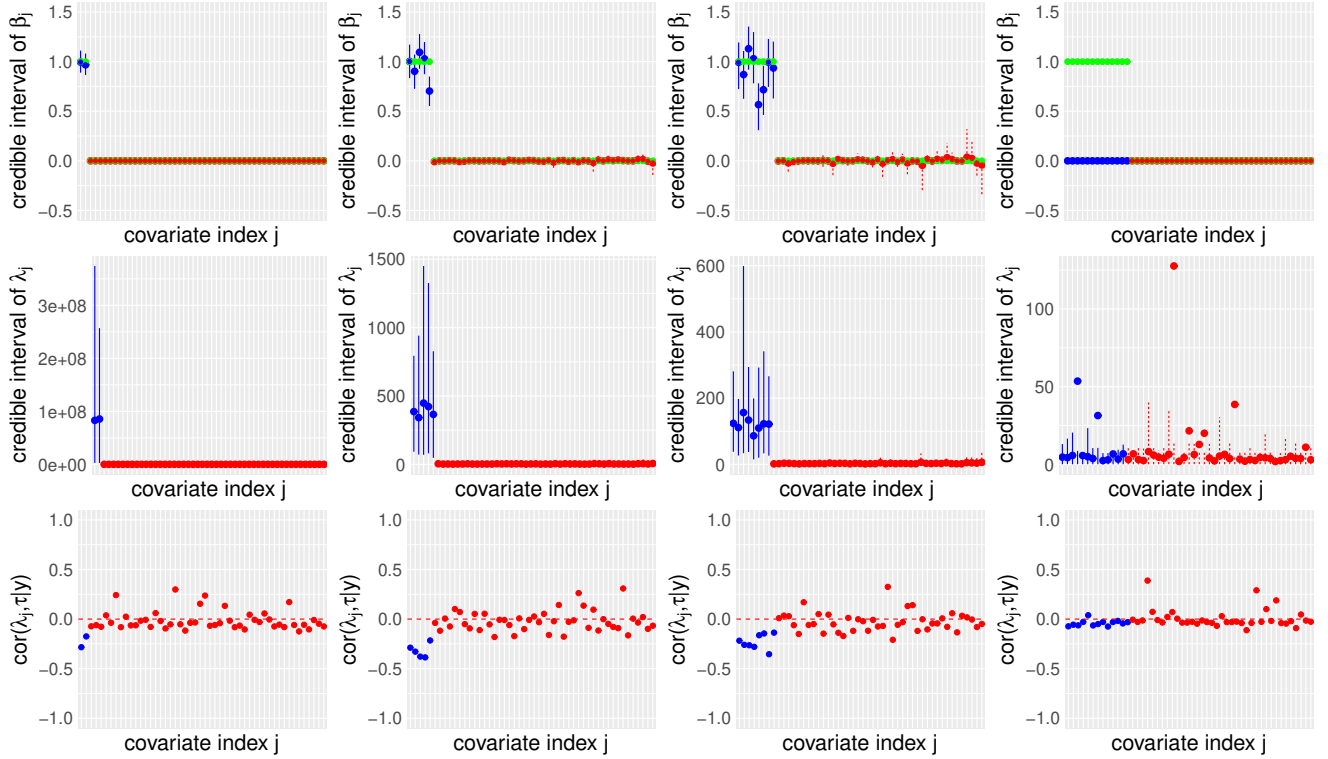


Figure 2: Results of posterior inference by using the Horseshoe under varying sparsity levels: \mathcal{A}_1 (first column), \mathcal{A}_2 (second column), \mathcal{A}_3 (third column), and \mathcal{A}_4 (fourth column). The results of posterior inference corresponding to signals and noises are colored in blue and red, respectively, and the truth β_0 is colored in green. Posterior means of τ corresponding to the four datasets are $1.41 \cdot 10^{-6}$ (\mathcal{A}_1), 0.05 (\mathcal{A}_2), 0.13 (\mathcal{A}_3), and $6.53 \cdot 10^{-15}$ (\mathcal{A}_4), respectively.

the first 50 coefficients of β are plotted. The results of posterior inference corresponding to signals and noises are colored blue and red, respectively, and the true coefficient vector β_0 is colored green. The posterior means of τ corresponding to the four datasets are $1.41 \cdot 10^{-6}$ (\mathcal{A}_1), 0.05 (\mathcal{A}_2), 0.13 (\mathcal{A}_3), and $6.53 \cdot 10^{-15}$ (\mathcal{A}_4), respectively. Hence, the posterior mean of τ gradually increases as the sparsity level increases and after some threshold it suddenly drops to a very small number. The relationship between the local $\{\lambda_j\}_{j=1}^p$ and the global τ scale parameters is the key to comprehend how the Horseshoe detects signals from a posteriori perspective. Observe that τ is associated with the sparsity level [37], and is expected to be large in presence of a relatively high number of signals. The panels on the third row of the Figure 2 shows weak negative posterior correlation between λ_j and τ , $\text{cor}(\lambda_j, \tau|y)$, for each $j = 1, \dots, p$. As seen on the panels on the first and third rows of the Figure 2, the selected signals among the p coefficients $\{\beta_j\}_{j=1}^p$, saying $\{\beta_j\}_{j \in \mathcal{Q}}$,

$\mathcal{Q} \subset \mathcal{P} = \{1, \dots, p\}$, are those whose corresponding posterior correlations $\{\text{cor}(\lambda_j, \tau|y)\}_{j \in \mathcal{Q}}$ attain even stronger negative values than others $\{\text{cor}(\lambda_j, \tau|y)\}_{j \in \mathcal{P} - \mathcal{Q}}$. This implies that if there are no discriminable differences among the correlations $\{\text{cor}(\lambda_j, \tau|y)\}_{j \in \mathcal{P}}$ then the Horseshoe loses its signal detection mechanism and a collapse takes place, as seen on the panel of the fourth column.

It is important to emphasize that the collapsing behavior for the Horseshoe was pointed out by several authors but did not draw much attention in the literature. Recently, [52] in discussion of [48] discussed dangers of collapsing of marginal maximum-likelihood estimator for the global-scale parameter when the sparsity level is very small. On the other hand, the focal point of our research is to investigate collapsing behavior of the fully Bayesian Horseshoe estimator, i.e., where τ is assumed to follow $\mathcal{C}^+(0, 1)$ in the case when the sparsity level is moderately large.

4.1 Restricted tail-heaviness of Horseshoe

In the following, we shall aim to understand the tail-heaviness of the Horseshoe and how it may affect the signal detection mechanism. Suppose $\sigma^2 = 1$ and consider the univariate covariate free formulation of (1) as $y|\beta \sim \mathcal{N}_1(\beta, 1)$, $\beta|\lambda, \tau \sim \mathcal{N}_1(0, \lambda^2\tau^2)$, and $\lambda \sim \mathcal{C}^+(0, 1)$ with fixed $\tau > 0$. Our focus is on the tail part of the marginal density of the Horseshoe conditional on τ , as denoted by $\pi_{\text{HS}}(\beta|\tau)$ in (1) of the supplementary document. We also aim to understand the importance of tail-heaviness to handle moderately larger number of signals and the behavior of the tail of $\pi_{\text{HS}}(\beta|\tau)$ as τ changes. To answer this question, we begin with defining tail-heaviness of a measurable function ρ using the notion of regular variation [17, 20, 27, 28].

Definition 1. *A measurable function $\rho : (0, \infty) \rightarrow (0, \infty)$ is called regularly varying of index α , if there exists $\alpha \in \mathfrak{R}$ such that $\lim_{x \rightarrow \infty} \rho(cx)/\rho(x) = c^{-\alpha}$, for any $c > 0$. If $\alpha = 0$, then ρ is said to be slowly varying.*

In general, every regularly varying function ρ of index α has a representation $\rho(x) = L(x) \cdot x^{-\alpha}$ where L is a slowly varying function [27]. In extreme value theory, the Definition 1 is utilized to quantify the tail-heaviness of a positive random variable $X \sim F$ where F is the distribution function of X by replacing the measurable function ρ in the Definition 1 with the tail (survival)

function $\bar{F} = 1 - F$. This leads to $\bar{F}(x) = L(x) \cdot x^{-\alpha}$, where α represents the decay rate \bar{F} at infinity, called the tail-index of the random variable X or the tail-index of the density $f = F'$ [13, 17]. The reciprocal $\xi = 1/\alpha$ is called the shape parameter. This notion can be generalized to a random variable supported on $(-\infty, \infty)$ in a similar fashion by quantifying the tail behavior both at ∞ and $-\infty$.

Proposition 2. *Assume $\beta|\lambda, \tau \sim \mathcal{N}_1(0, \tau^2\lambda^2)$, $\lambda \sim C^+(0, 1)$, and $\tau > 0$. Then the tail-index of $\pi_{\text{HS}}(\beta|\tau)$ is $\alpha = 1$ for any $\tau > 0$.*

Proposition 2 is proved in §7 of the supplementary materials. The tail-index of half-Cauchy density $\mathcal{C}^+(0, 1)$ is $\alpha = 1$ [17]. Proposition 2 implies that the tail-heaviness of the marginal density $\pi_{\text{HS}}(\beta|\tau)$ inherits that of the local-scale density $\pi(\lambda) = \mathcal{C}^+(0, 1)$, and is fixed for any choice of the global-scale parameter $\tau > 0$. Hence the Horseshoe is unable to adjust the tail of $\pi_{\text{HS}}(\beta|\tau)$ to deal with various sparsity regimes [36]. Horseshoe is designed to perform well in a very sparse situation since setting $\alpha = 1$ is sufficient to put an enough mass on the tail region of $\pi_{\text{HS}}(\beta|\tau)$. This is also backed up by the numerical results, refer to left panels in the Figure 2. However, as the sparsity level increases, the requirement of placing more mass in the tail part of $\pi_{\text{HS}}(\beta|\tau)$ increases and $\alpha = 1$ is not sufficiently large to achieve this. This is shown in the right panel of Figure 2.

5. A REMEDY FOR THE COLLAPSING BEHAVIOR

We propose a fully Bayesian remedy for the collapsing behavior of the Horseshoe by introducing a new hierarchical formulation. In the following, we expand the existing framework of global-local shrinkage priors to a new framework called the global-local-tail shrinkage priors. The global-scale parameter τ is introduced as the scale parameter of a local-scale density f with tail-index α , or equivalently, shape parameter $\xi = 1/\alpha$.

$$\beta_j|\lambda_j, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2\sigma^2), \quad \sigma^2 \sim h(\sigma^2), \quad (j = 1, \dots, p), \quad (6)$$

$$\lambda_j|\tau, \xi \sim f(\lambda_j|\tau, \xi), \quad (j = 1, \dots, p), \quad (7)$$

$$(\tau, \xi) \sim g(\tau, \xi), \quad (8)$$

where f is a density supported on $(0, \infty)$ with the scale parameter $\tau > 0$ and the shape parameter $\xi > 0$ such that the distribution function $F(\lambda|\tau, \xi) = \int_0^\lambda f(t|\tau, \xi)dt$ for $\lambda > 0$ satisfies $\bar{F}(\lambda|\tau, \xi) = 1 - F(\lambda|\tau, \xi) = L(\lambda) \cdot \lambda^{-(1/\xi)}$ such that L is a slowly varying function. h is a density supported on $(0, \infty)$ and g is a joint density supported on $(0, \infty) \times (0, \infty)$. We call the formulation (6) – (8) the global-local-tail shrinkage prior because the presence of the shape parameter ξ which measures the tail-heaviness of f .

Table 1 lists examples of f with unit scale parameter $\tau = 1$. All distributions in the table are supported on $(0, \infty)$ with a positive shape parameter $\xi > 0$. The half-Cauchy distribution and the half-Levy distributions are derived from the half- α -stable distribution with the tail-index α by fixing α to be 1 and 1/2, respectively. More examples for f can be found in [17, 23].

The Horseshoe is a member of the global-local-tail shrinkage priors (6) – (8) when (i) the local-scale density f is chosen by the half- α -stable density, and (ii) the tail-index α is fixed with 1 or equivalently $\xi = 1$.

Table 1: Unit scaled densities for f in (7)

	$f(\lambda \tau = 1, \xi)$	Shape parameter ξ
Half- α -stable distribution	non-closed form	ξ
Half-Cauchy distribution	$2\{\pi(1 + \lambda^2)\}^{-1}$	1
Half-Levy distribution	$\lambda^{-3/2}\exp\{-1/(2\lambda)\}/\sqrt{2\pi}$	2
Loggamma distribution	$\{(1 + \lambda)^{-(1/\xi+1)}\}/\xi$	ξ
Generalized extreme value distribution	$\exp\{-(1 + \xi\lambda)^{-1/\xi}\}(1 + \xi\lambda)^{-(1/\xi+1)}$	ξ
Generalized Pareto distribution	$(1 + \xi\lambda)^{-(1/\xi+1)}$	ξ

In the formulation of global-local-tail shrinkage prior (6)-(8), we use the generalized Pareto distribution (GPD) [35], $\pi(x) = \mathcal{GPD}(x|\tau, \xi) = (1/\tau) \cdot (1 + \xi x/\tau)^{-(1/\xi+1)}$ for the local-scale density f , and a truncated inverse-gamma-lognormal joint density for the joint density g , $g(\tau, \xi) = \mathcal{IG}(\tau|p/\xi+1, 1)\mathcal{I}_{(0,\infty)}(\tau) \cdot \{\log \mathcal{N}(\xi|\mu, \rho^2)\mathcal{I}_{(1/2,\infty)}(\xi)\}/D$, where $D = D(\mu, \rho^2) = \int_{1/2}^\infty \log \mathcal{N}(\xi|\mu, \rho^2)d\xi$




Figure 3: DAG representation of $y|\beta, \sigma^2 \sim \mathcal{N}_n(X\beta, \sigma^2 I_n)$ and $\beta \sim \pi_{\text{GLT}}(\beta)$.

is the normaliser of $g(\tau, \xi)$,

$$\beta_j|\lambda_j, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2 \sigma^2), \quad \sigma^2 \sim \pi(\sigma^2) \propto 1/\sigma^2, \quad (j = 1, \dots, p), \quad (9)$$

$$\lambda_j|\tau, \xi \sim \mathcal{GPD}(\tau, \xi), \quad (j = 1, \dots, p), \quad (10)$$

$$\tau|\xi \sim \mathcal{IG}(p/\xi + 1, 1), \quad (11)$$

$$\xi \sim \log \mathcal{N}(\mu, \rho^2) \mathcal{I}_{(1/2, \infty)}, \quad \mu \in \mathfrak{R}, \rho^2 > 0. \quad (12)$$

We call this specific hierarchical form (9)-(12) the GLT prior, denoted as $\beta \sim \pi_{\text{GLT}}(\beta)$. A directed asymmetric graphical (DAG) representation of $y|\beta, \sigma^2 \sim \mathcal{N}_n(X\beta, \sigma^2 I_n)$ and $\beta \sim \pi_{\text{GLT}}(\beta)$ is shown in Figure 3. A full description of posterior computation is provided in §5 of the supplementary document. §5 contains an automatic-tuning for the hyper-parameter, μ and ρ^2 , using extreme value theory which enables learning of ξ adaptive to the unknown sparsity level. Henceforth, we omit to write the suffix ‘GLT’ except for $\pi_{\text{GLT}}(\beta)$, and to avoid notational confusion with the results from the Horseshoe the suffix ‘HS’ is explicitly used as we did in Lemma 1 in the supplementary materials and Proposition 2.

Figure 4 displays the results of posterior inference obtained using the GLT prior when applied to the same four artificial datasets \mathcal{A}_l ($l = 1, 2, 3, 4$) used in Figure 2. The posterior correlations $\{\text{cor}(\lambda_j, \xi|y)\}_{j=1}^p$ are additionally plotted in the panels on the fourth row of the Figure 4. Note that the GLT prior can detect signals in the dataset \mathcal{A}_4 unlike the Horseshoe. The posterior means of global-scale parameter τ corresponding to the four datasets are 0.003 (\mathcal{A}_1), 0.003 (\mathcal{A}_2), 0.004 (\mathcal{A}_3), and 0.004 (\mathcal{A}_4), respectively. The posterior means of the shape parameter ξ corresponding to the four datasets are 2.010 (\mathcal{A}_1), 2.134 (\mathcal{A}_2), 2.235 (\mathcal{A}_3), 2.347 (\mathcal{A}_4), respectively. The signal

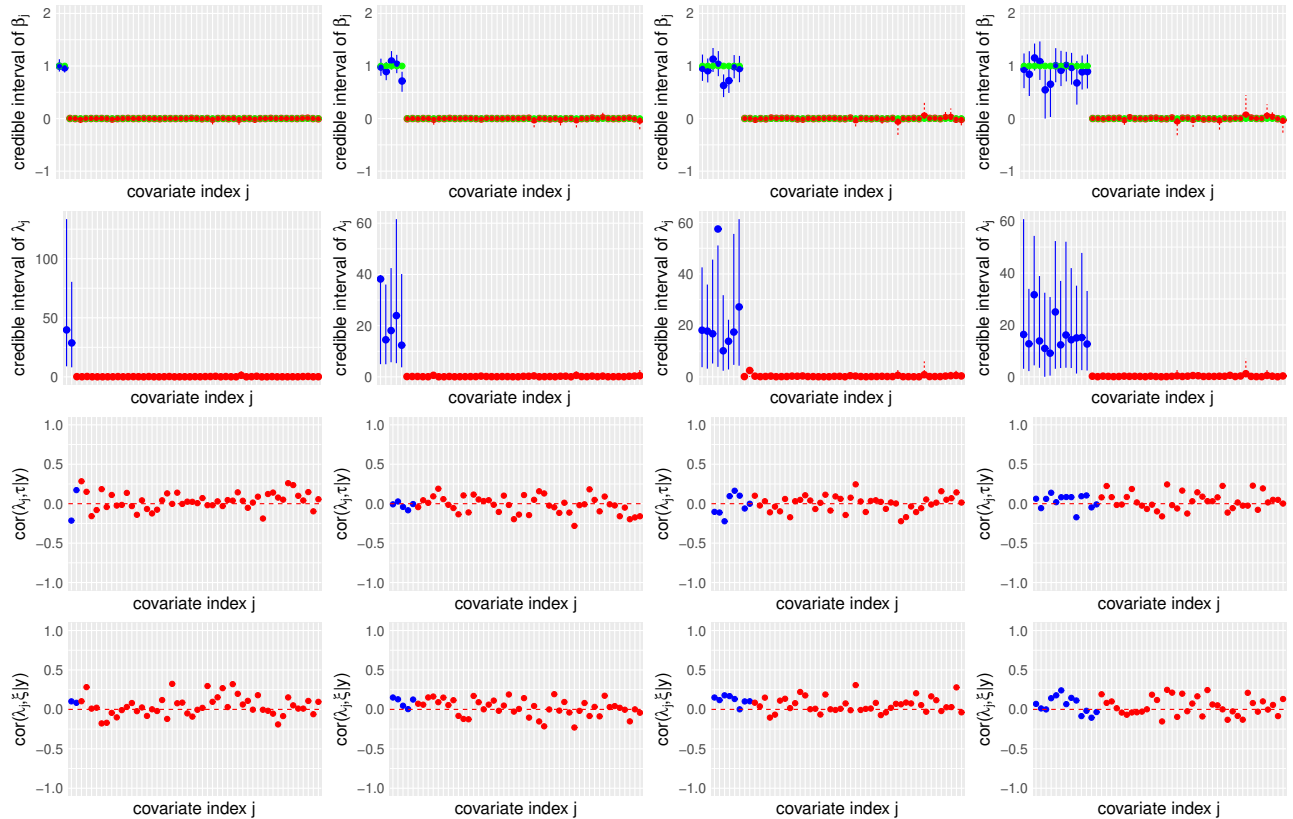


Figure 4: Results of posterior inference obtained using the GLT prior when applied to the same four artificial datasets used in the Figure 2: \mathcal{A}_1 (first column), \mathcal{A}_2 (second column), \mathcal{A}_3 (third column), and \mathcal{A}_4 (fourth column). Posterior means of (τ, ξ) corresponding to the four datasets are $(0.003, 2.010)$ (\mathcal{A}_1), $(0.004, 2.134)$ (\mathcal{A}_2), $(0.004, 2.235)$ (\mathcal{A}_3), and $(0.004, 2.347)$ (\mathcal{A}_4), respectively.

detection mechanism of the GLT prior is accomplished by the p local-scale parameters $\{\lambda_j\}_{j=1}^p$ which embed the uncertainty about both the global-scale parameter τ and the shape parameter ξ . Also, since the estimates of τ and ξ are nearly independent from those of $\{\lambda_j\}_{j=1}^p$, refer to the third and the fourth rows of the Figure 4. The estimates of τ are not susceptible to the collapse even when the number of true signals is relatively large.

We must point out at this point that our strategy of introducing a tail controlling mechanism is not a unique remedy. We have implemented simulation studies for variant versions for the Horseshoe, namely truncated Horseshoe [44], Horseshoe-plus [3], and regularized Horseshoe [36], under the three scenarios discussed in Section 6, and we observed a similar collapse for the truncated Horseshoe but not for the other two variants. In majority of the cases, we observe superior performance of the the GLT prior and the truncated Horseshoe compared to the others, with the GLT prior achieve better mean squared error for the signal estimation part. Refer to §4 of the supplementary materials.

5.1 Properties of the GLT prior

For the purpose of prior analysis for the GLT prior, we shall work with the univariate form, $y|\beta \sim \mathcal{N}_1(\beta, 1)$, $\beta|\lambda \sim \mathcal{N}_1(0, \lambda^2)$, and $\lambda|\tau, \xi \sim \mathcal{GPD}(\tau, \xi)$, with fixed $\tau > 0$ and $\xi > 1/2$.

Proposition 3 (Marginal density of the GLT prior).

(a) Suppose $\beta|\lambda \sim \mathcal{N}_1(0, \lambda^2)$, $\lambda \sim \mathcal{GPD}(\tau, \xi)$, $\tau > 0$ and $\xi > 1/2$. Then:

$$\pi(\beta|\tau, \xi) = \sum_{k=0}^{\infty} a_k \{\psi_k^S(\beta) + \psi_k^R(\beta)\}, \quad (13)$$

where $K = 1/(\tau 2^{3/2} \pi^{1/2})$, $Z(\beta) = \beta^2 \xi^2 / (2\tau^2)$, $a_k = (-1)^k \cdot K \cdot \binom{1/\xi+k}{k}$, $\psi_k^S(\beta) = E_{k/2+1}\{Z(\beta)\}$, and $\psi_k^R(\beta) = Z(\beta)^{-\frac{1+1/\xi+k}{2}} \gamma\{(1 + 1/\xi + k)/2, Z(\beta)\}$.

Two special functions are used in (13): (i) the generalized exponential-integral function of real order [12, 30] $E_s(x) = \int_1^\infty e^{-xt} t^{-s} dt$ ($x > 0$, $s \in \mathfrak{R}$), and (ii) the incomplete lower gamma function $\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt$ ($s, x \in \mathfrak{R}$). The generalized binomial coefficient $\binom{1/\xi+k}{k}$ is $(1/\xi + k)(1/\xi + k - 1) \cdots (1/\xi + 1)/k!$ if $k \in \{1, 2, \dots\}$, and zero if $k = 0$.

(b) Suppose $\lambda \sim \mathcal{GPD}(\tau, \xi)$, $\kappa = 1/(1 + \lambda^2) \in (0, 1)$, $\tau > 0$ and $\xi > 1/2$. Then:

$$\pi(\kappa|\tau, \xi) = \frac{\tau^{1/\xi}}{2} \cdot \frac{\kappa^{1/(2\xi)-1}(1-\kappa)^{-1/2}}{\{\tau\kappa^{1/2} + \xi(1-\kappa)^{1/2}\}^{(1+1/\xi)}}. \quad (14)$$

Proposition 3 is proved in §8 of the supplementary materials. Figure 5 displays the marginal densities of β obtained from the Horseshoe $\pi_{\text{HS}}(\beta|\tau)$ ($\tau > 0$) (1) in the supplementary materials and the GLT prior $\pi(\beta|\tau, \xi)$ ($\tau > 0, \xi > 1/2$) (13) for different values of values τ and ξ . Note that $\pi(\beta|\tau, \xi)$ (13) is analytically expressed as an alternating series whose summands are separated into two terms: $\{\psi_k^{\text{S}}(\beta)\}_{k=0}^{\infty}$ and $\{\psi_k^{\text{R}}(\beta)\}_{k=0}^{\infty}$. The superscripts S and R stand for shrinkage and robustness, respectively. Corollary 4 implies that the roles of $\{\psi_k^{\text{S}}(\beta)\}_{k=0}^{\infty}$ and $\{\psi_k^{\text{R}}(\beta)\}_{k=0}^{\infty}$ in the GLT prior are separably interpretable in the limiting sense as $|\beta| \rightarrow 0$ and $|\beta| \rightarrow \infty$.

Corollary 4. Suppose $\beta|\lambda \sim N_1(0, \lambda^2)$, $\lambda \sim \mathcal{GPD}(\tau, \xi)$, $\tau > 0$, and $\xi > 1/2$. Let $k \in \{0\} \cup \{1, 2, \dots\}$. Then:

(a) If $k = 0$, then $\lim_{|\beta| \rightarrow 0} \psi_k^{\text{S}}(\beta) = \infty$; if $k \in \{1, 2, \dots\}$, then $\lim_{|\beta| \rightarrow 0} \pi_k^{\text{S}}(\beta) = 2/k < \infty$.

(b) If $k \in \{0\} \cup \{1, 2, \dots\}$, then $\lim_{|\beta| \rightarrow \infty} \psi_k^{\text{S}}(\beta) = 0$ with squared exponential rate.

(c) If $k \in \{0\} \cup \{1, 2, \dots\}$, then $\lim_{|\beta| \rightarrow 0} \psi_k^{\text{R}}(\beta) = 2/(1 + 1/\xi + k) < \infty$.

(d) If $k \in \{0\} \cup \{1, 2, \dots\}$, then $\psi_k^{\text{R}}(\beta)$ is regularly varying with index $1 + 1/\xi + k$.

Interpretations of the Corollary 4 are as follows. (a) implies that the marginal of β of the GLT prior $\pi(\beta|\tau, \xi)$ possesses infinite spike at origin for any $\tau > 0, \xi > 1/2$, as seen in the Figure 5. This is a feature shared by the Horseshoe [9] as well. The pole at zero is caused by the exponential-integral function of the first order: $\lim_{x \rightarrow 0^+} E_1(x) = \infty$ [12] and generates a very strong shrinkage on β towards zero. By (a) and (c) from the Corollary 4, we have $\lim_{|\beta| \rightarrow 0} \psi_k^{\text{R}}(\beta) = 2/(1 + 1/\xi + k) < \lim_{|\beta| \rightarrow 0} \pi_k^{\text{S}}(\beta) = 2/k$, $k \in \{1, 2, \dots\}$, which implies that the contribution of $\{\pi_k^{\text{S}}(\beta)\}_{k=0}^{\infty}$ is more than that of $\{\psi_k^{\text{R}}(\beta)\}_{k=0}^{\infty}$ in the shrinkage of the β . Squared exponential decay rates of the terms in $\{\pi_k^{\text{S}}(\beta)\}_{k=0}^{\infty}$ at $|\beta| \rightarrow \infty$ in (b) implies that the contribution of $\{\pi_k^{\text{S}}(\beta)\}_{k=0}^{\infty}$ in controlling the tail of $\pi(\beta|\tau, \xi)$ is negligible as $|\beta|$ goes to infinity. Finally, (d) implies the

marginal prior $\pi(\beta|\tau, \xi)$ has a systematic mechanism to control the tails, $\{\psi_k^R(\beta)\}_{k=0}^\infty$, achieved by controlling ξ .

We call the sequence of functions $\{\psi_k^R(\beta)\}_{k=0}^\infty$ tail lifters as their main roles are to lift the tail part of the density $\pi(\beta|\tau, \xi)$ by increasing ξ . The presence of tail lifters in the marginal prior $\pi(\beta|\tau, \xi)$ provides a great flexibility to the shape of the density as shown in the panels in Figure 5. This is particularly useful to handle various sparsity regimes.

In contrast, $\pi_{\text{HS}}(\beta|\tau)$ ((1) in the supplementary material) does not have any tail controlling mechanism (refer to Corollary 2). This is particularly problematic when τ is estimated to be very small (say $\tau = 0.001$). Panels in the second row in Figure 5 show a mismatch between the theoretical support \mathfrak{R} and numerical support $(-\epsilon, \epsilon)$, $\epsilon \approx 0$ of the marginal density $\pi_{\text{HS}}(\beta|\tau = 0.001)$ (1). If τ is extremely small, say, $\tau = 10^{-10}$ (a typical phenomenon in the moderately sparse regime), then the density $\pi_{\text{HS}}(\beta|\tau = 10^{-10})$ is numerically approximated by the Dirac-delta function $\delta_0(\beta)$ causing the collapse.

Corollary 5. *Let $\lambda|\tau, \xi \sim \mathcal{GPD}(\tau, \xi)$, $\kappa = 1/(1 + \lambda^2) \in (0, 1)$, $\tau > 0$, and $\xi > 1/2$. Then:*

(a) $\lim_{\kappa \rightarrow 1^-} \pi(\kappa|\tau, \xi) = \infty$ and $\lim_{\kappa \rightarrow 0^+} \pi(\kappa|\tau, \xi) = \infty$.

(b) $\pi(\kappa|\tau = 1, \xi = 1) = \{\kappa^{-1/2}(1 - \kappa)^{-1/2}\} / [2 \cdot \{\kappa^{1/2} + (1 - \kappa)^{1/2}\}^2]$.

In the univariate case, the posterior mean of β can be represented as $E[\beta|y] = (1 - E[\kappa|y])y$ with $\kappa = 1/(1 + \lambda^2) \in (0, 1)$. Probabilities of the regions $(1 - \epsilon, 1)$ and $(0, \epsilon)$, $\epsilon \approx 0$ under $\pi(\kappa|\tau, \xi)$ are related with the shrinkage and the robustness [9]. The infinite spikes of $\pi(\kappa|\tau, \xi)$ at $k = 0$ and $k = 1$ imply that the GLT prior has the desired shrinkage property. The density $\pi(\kappa|\tau = 1, \xi = 1)$ is not standard, but resembles a ‘horseshoe’.

Figure 6 displays two densities $\pi_{\text{HS}}(\kappa|\tau)$ in (2) of the supplementary materials and $\pi(\kappa|\tau, \xi)$ (14), with different values of τ and ξ . When $\tau = 1$, the top panels demonstrate Horseshoe-like shapes for both $\pi_{\text{HS}}(\kappa|\tau = 1)$ and $\pi(\kappa|\tau = 1, \xi)$. However, when $\tau = 0.001$, the apparent difference is shown on the bottom-middle panel, where $\pi_{\text{HS}}(\kappa|\tau = 0.001)$ places essentially zero-mass on $(0, \epsilon)$, $\epsilon \approx 0$. This implies that the robustness property of the Horseshoe is deteriorated

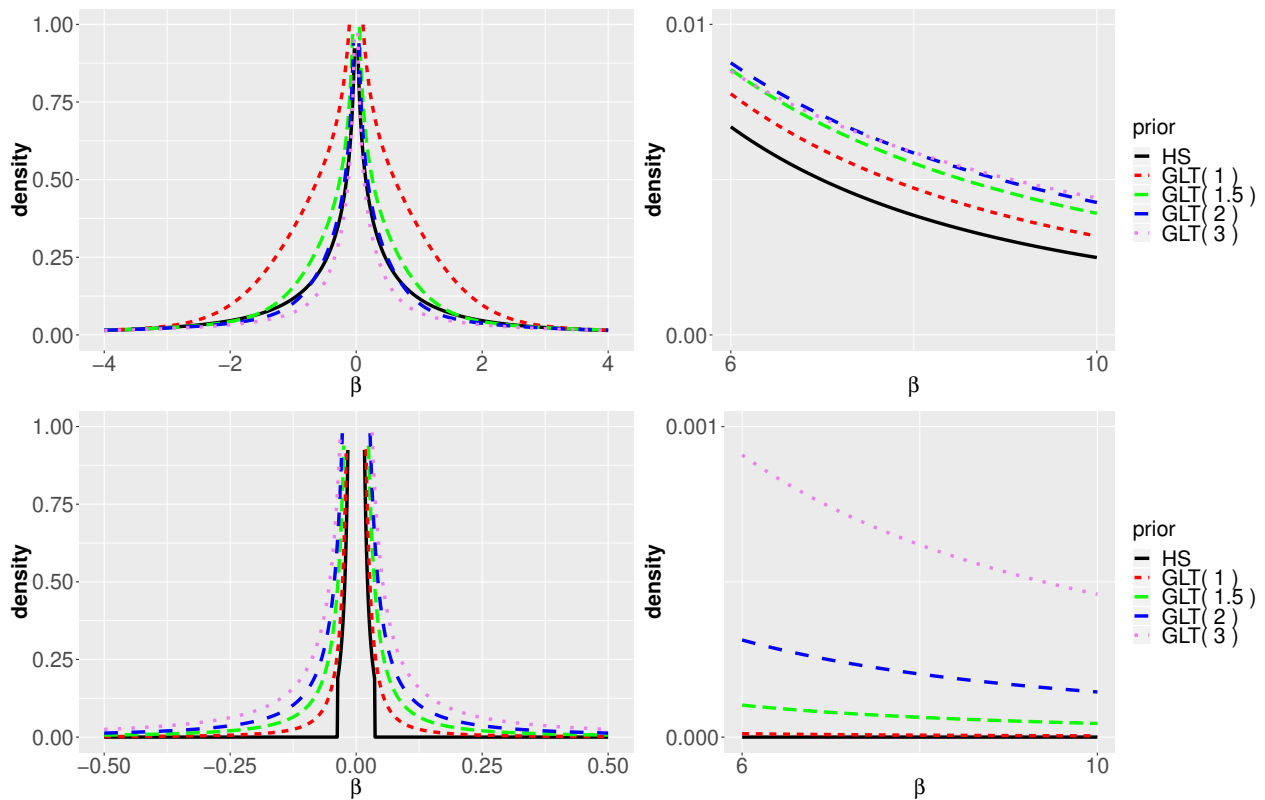


Figure 5: Comparison between two densities, $\pi_{\text{HS}}(\beta|\tau)$ ((a) in Lemma 1 in the supplementary material) and $\pi(\beta|\tau, \xi)$ ((a) in Proposition 3): $\tau = 1$ (top panels) and $\tau = 0.001$ (bottom panels). The density $\pi_{\text{HS}}(\beta|\tau)$ is colored in black, and densities $\pi(\beta|\tau, \xi)$ are colored in red ($\xi = 1$), green ($\xi = 1.5$), blue ($\xi = 2$), and violet ($\xi = 3$), respectively.

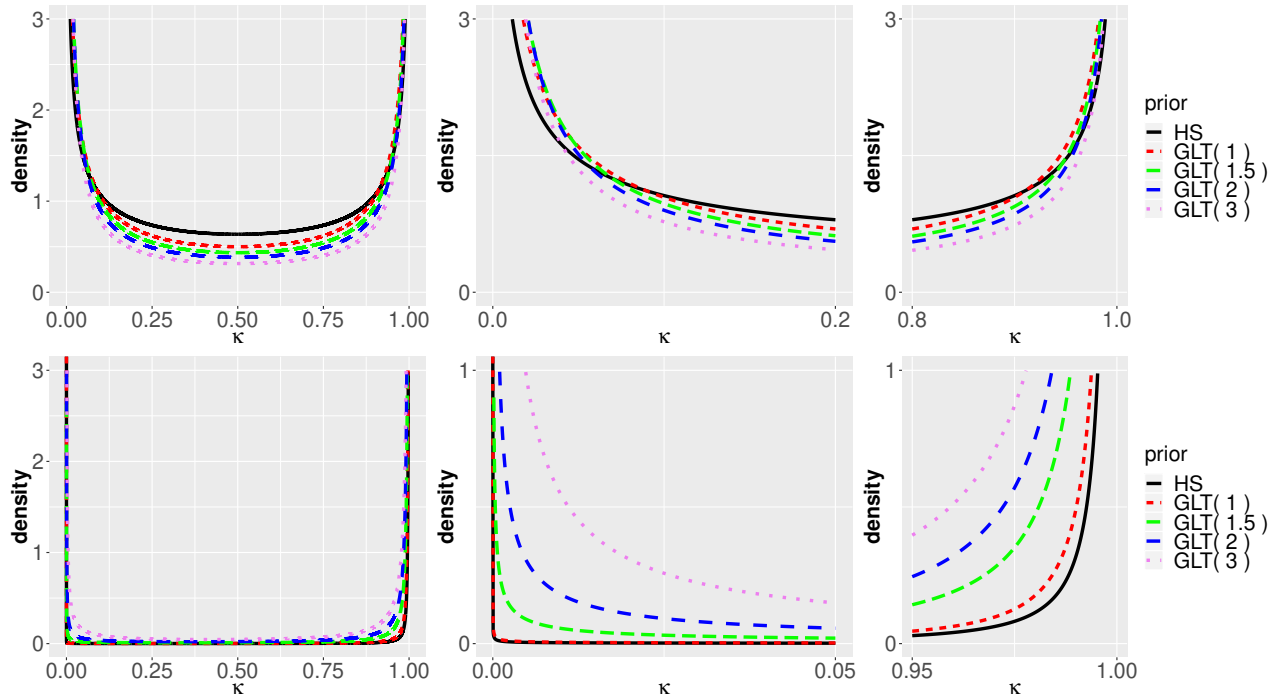


Figure 6: Comparison between two densities of random shrinkage weight κ , $\pi_{\text{HS}}(\kappa|\tau)$ ((b) in Lemma 1 in the supplementary material) and $\pi(\kappa|\tau, \xi)$ ((b) in Proposition 3): $\tau = 1$ (top panels) and $\tau = 0.001$ (bottom panels). The density $\pi_{\text{HS}}(\kappa|\tau)$ is colored in black, and densities $\pi(\kappa|\tau, \xi)$ are colored in red ($\xi = 1$), green ($\xi = 1.5$), blue ($\xi = 2$), and violet ($\xi = 3$), respectively.

when τ is very small. On the other hand, the GLT prior $\pi(\kappa|\tau = 0.001, \xi)$ places positive mass on $(0, \epsilon)$, $\epsilon \approx 0$, and the mass increases as ξ increases. This implies that the robustness property of the GLT prior is maintained even when τ is very small, and is adjustable by controlling ξ .

6. SIMULATION STUDY

In §4 we discussed how to generate an artificial high-dimensional data $(y, X) \in \mathfrak{R}^n \times \mathfrak{R}^{n \times p}$ corresponding to a simulation settings $(n, p, q, \varrho, \text{SNR})$ from a sparse linear regression (1) when the truth $\beta_0 = (\beta_{0,1}, \dots, \beta_{0,q}, \beta_{0,q+1}, \dots, \beta_{0,p})^T \in \mathfrak{R}^p$ is specified by $\beta_{0,1} = \dots = \beta_{0,q} = 1$ and $\beta_{0,q+1} = \dots = \beta_{0,p} = 0$. Now, we conduct a replicated study to compare the performance of the Horseshoe and the GLT prior as follows. First, set $(n, p) \in \{(100, 500), (200, 1000)\}$, $s = q/p = 0.01$, $\text{SNR} = 5$, and $\varrho = 0$, and then separately consider the following three scenarios by varying one environmental value while fixing others; 1) varied sparsity level q/p from 0.001 to 0.1, 2) varied ϱ from 0 to 0.5 and 3) varied SNR from 2 to 10.

We separately report the medians of mean squared errors (MSE) corresponding to signal and noise coefficients measured across the 50 replicated datasets. Let $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T \in \mathfrak{R}^p$ is the posterior mean obtained by using either the Horseshoe or the GLT prior: then, MSE corresponding to signals and noises are defined by $\text{MSE}_S = (1/q) \sum_{j=1}^q (\hat{\beta}_j - 1)^2$, $\text{MSE}_N = \{1/(p - q)\} \sum_{j=q+1}^p (\hat{\beta}_j)^2$. Note that when collapse takes place, the posterior mean $\hat{\beta}$ numerically becomes the p -dimensional zero vector. Hence, the two metrics MSE_S and MSE_N numerically become 1 and 0, respectively in that case. The posterior computations for the Horseshoe and the GLT prior are fully automated and tuning procedures are not required, hence, a fair comparison can be achieved based on the metrics.

Figure 7 displays the medians of MSE_S , MSE_N , and posterior means of τ and ξ under Scenario 1. The top and bottom panels correspond to $(n, p) = (100, 500)$ and $(n, p) = (200, 1000)$, respectively. To be specific, the top panel corresponds to the setting $(n = 100, p = 500, q, \varrho = 0, \text{SNR} = 5)$ with $q \in \{1, 6, 11, 16, 22, 27, 32, 37, 43, 48\}$, and the bottom panel corresponds to $(n = 200, p = 1000, q, \varrho = 0, \text{SNR} = 5)$ with $q \in \{1, 11, 22, 32, 43, 53, 64, 74, 85, 95\}$ so that the sparsity level q/p varies from 0.001 to 0.1.

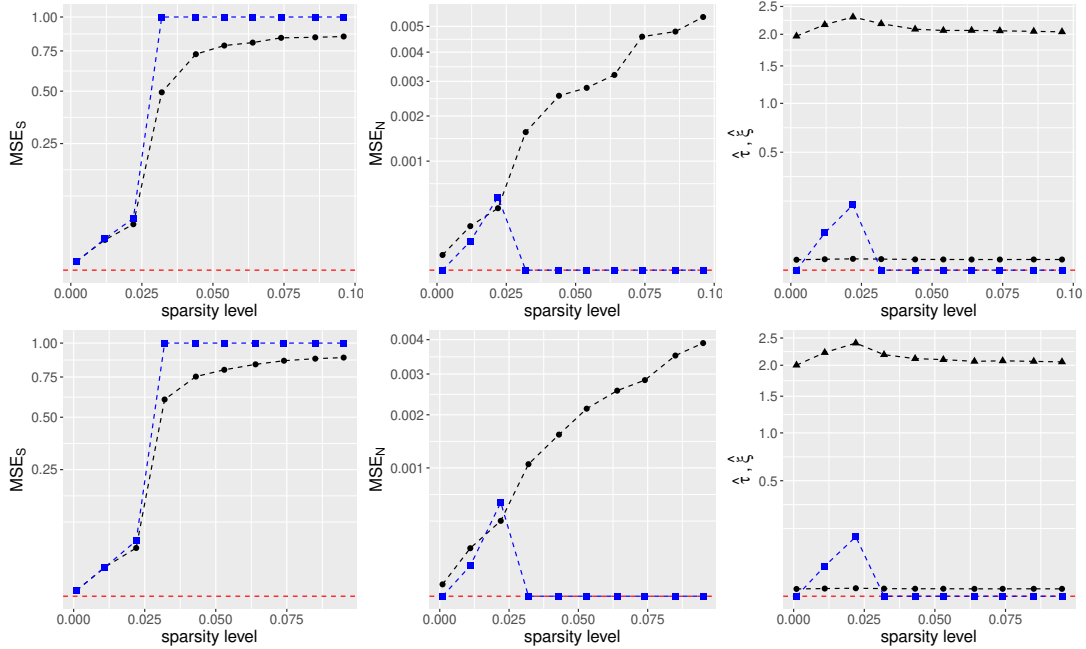


Figure 7: Medians of MSE_S , MSE_N , and posterior means of τ and ξ across different sparsity level q/p : $(n, p) = (100, 500)$ (top panel) and $(n, p) = (200, 1000)$ (bottom panel). Results from the GLT prior are marked with black circular dot \bullet ; black triangular dot (rightmost panel) represents the posterior mean of ξ . Results from the Horseshoe are marked with blue square dot. The red dotted horizontal line represents zeros.

The Horseshoe performs well if the sparsity level q/p is less than $11/500 = 22/1000 = 0.022$, but suddenly collapses beyond this due to a sharp decrease in the posterior mean of τ . On the other hand, the GLT prior does not collapse and the posterior means of the τ are maintained at around 0.004, and the posterior means of ξ increase as the sparsity level increases to $q/p = 0.022$.

Figure 8 displays the medians of MSE_S , MSE_N , and posterior means of τ and ξ under Scenario 2. MSE_S are MSE_N obtained from the both priors increase as ρ increases. The GLT prior shows better signal detection whereas the Horseshoe shows slightly better noise shrinkage. When $(n, p) = (200, 1000)$ and $\rho = 0.5$ the GLT prior outperforms the Horseshoe as seen on the bottom-middle panel.

Figure 9 displays the medians of MSE_S , MSE_N , and posterior means of τ and ξ under Scenario 3. MSE_S are MSE_N obtained from the both priors monotonically decreases as the value of SNR increases. Both priors show excellent performances in shrinkage noises as seen in the middle panels. When $SNR = 2$, the GLT prior shows better performance than the Horseshoe as seen in

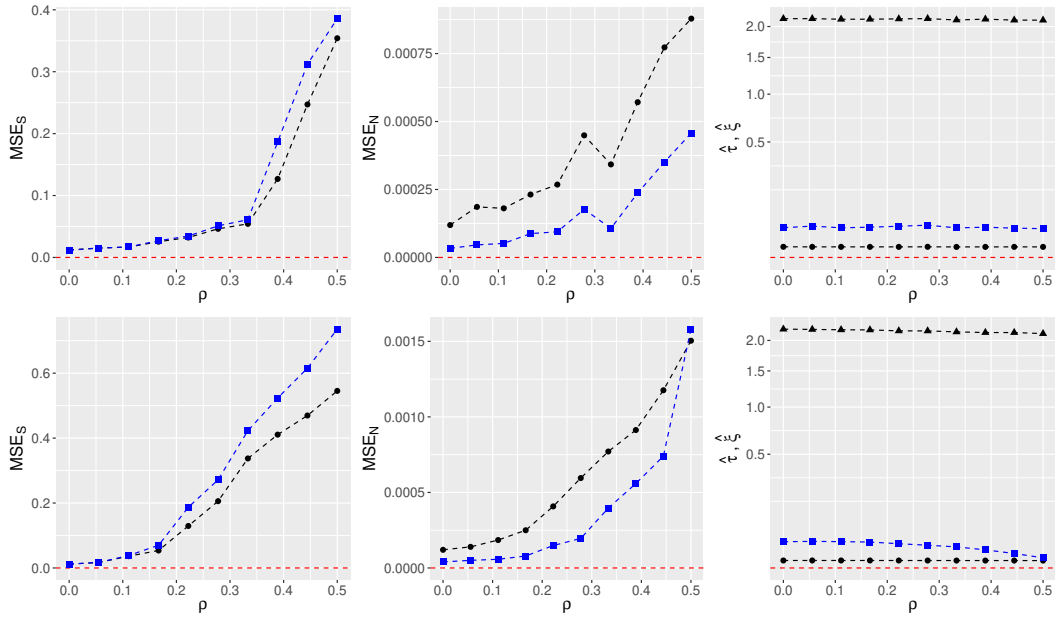


Figure 8: Medians of MSE_S , MSE_N , and posterior means of τ and ξ across different sparsity level q/p : $(n, p) = (100, 500)$ (top panel) and $(n, p) = (200, 1000)$ (bottom panel).

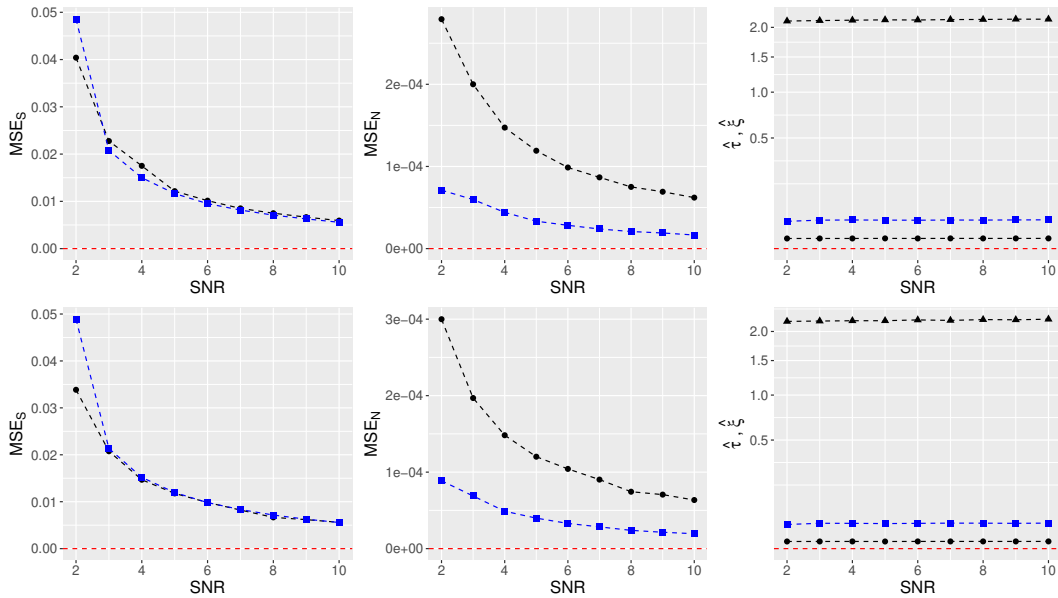


Figure 9: Medians of MSE_S , MSE_N , and posterior means of τ and ξ across different sparsity level q/p : $(n, p) = (100, 500)$ (top panel) and $(n, p) = (200, 1000)$ (bottom panel).

the left panels.

7. THE GLT PRIOR APPLIED TO THE PROSTATE CANCER DATA

To investigate the performance of the GLT prior on the prostate cancer data in Subsection 3.1 as the number of genes used increases, we constructed seven prostate datasets, $\mathcal{P}_1 = \{y_j\}_{j=1}^{p=50}$, $\mathcal{P}_2 = \{y_j\}_{j=1}^{p=100}$, $\mathcal{P}_3 = \{y_j\}_{j=1}^{p=200}$, $\mathcal{P}_4 = \{y_j\}_{j=1}^{p=500}$, $\mathcal{P}_5 = \{y_j\}_{j=1}^{p=1000}$, $\mathcal{P}_6 = \{y_j\}_{j=1}^{p=3000}$, and $\mathcal{P}_7 = \{y_j\}_{j=1}^{p=6033}$ such that $\mathcal{P}_1 \subset \mathcal{P}_2 \subset \mathcal{P}_3 \subset \mathcal{P}_4 \subset \mathcal{P}_5 \subset \mathcal{P}_6 \subset \mathcal{P}_7$, where \mathcal{P}_7 is the full dataset. In Subsection 3.1, we showed that the Horseshoe collapses when the number of genes used is 200 or more, as seen on the right panel of the Figure 1.

The results of posterior inference obtained using the GLT prior are shown on the Figure 10: the left and the right panels display the ordered pairs $\{(y_j, \hat{\beta}_j)\}_{j=1}^p$, when applied to the four datasets \mathcal{P}_l , $l = 1, 2, 3, 4$, and the three datasets \mathcal{P}_l , $l = 5, 6, 7$, respectively, where $\hat{\beta}_j$ is the posterior mean of β_j . The reversed- S -shape curves formed by pairs $\{(y_j, \hat{\beta}_j)\}_{j=1}^p$ for all datasets testify to the robustness property of GLT regardless of how many genes are used. This is primarily due to its automatic adaptation of the tail-heaviness to the data used. This is a clear advantage of GLT over the Horseshoe whose tail-heaviness is fixed; refer to Proposition 2.

The posterior means of the shape parameter ξ for the seven datasets are 1.620 (\mathcal{P}_1), 1.662 (\mathcal{P}_2), 1.789 (\mathcal{P}_3), 1.905 (\mathcal{P}_4), 1.991 (\mathcal{P}_5), 2.760 (\mathcal{P}_6), and 3.636 (\mathcal{P}_7), respectively. This monotonicity ξ is indicative of the fact that the GLT prior automatically adapts its tail-heaviness to accommodate more interesting genes as the number of genes used increases.

8. DISCUSSION

The purpose of this article is not to criticize the existing continuous shrinkage priors or Horseshoe in particular, but is simply to recognize that these priors are devised to produce meaningful results only in the regime where there are a handful of true signals, the so-called ultra-sparse regime. However, accompanied by the advancement in modern microarray technique and gene discovery, it is necessary to devise a sparsity inducing prior which works reasonably well across diverse sparsity regimes. The proposed GLT prior aims to address this gap in the literature. We also explored

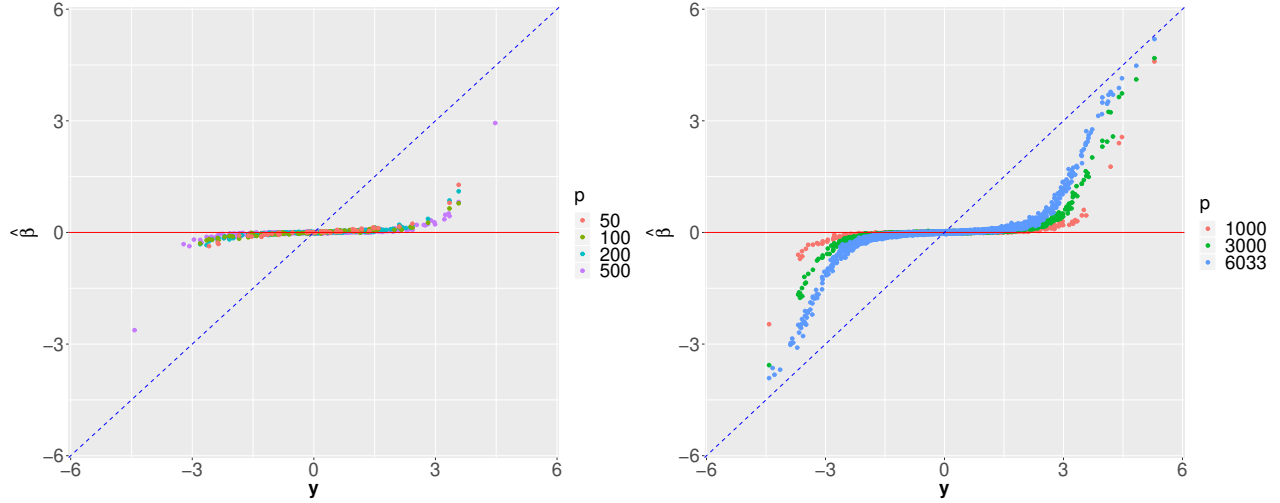


Figure 10: Posterior inference results obtained by the GLT prior applied seven prostate cancer datasets \mathcal{P}_l , $l = 1, \dots, 7$. Posterior means of (τ, ξ) corresponding to the seven datasets are $(0.0303, 1.620)$ (\mathcal{P}_1), $(0.0154, 1.662)$ (\mathcal{P}_2), $(0.0090, 1.789)$ (\mathcal{P}_3), $(0.0037, 1.905)$ (\mathcal{P}_4), $(0.0019, 1.991)$ (\mathcal{P}_5), $(0.0013, 2.760)$ (\mathcal{P}_6), and $(0.0013, 3.636)$ (\mathcal{P}_7), respectively.

application of the proposed GLT prior to a curve fitting study; refer to §2 of the supplementary materials.

We emphasize that delicate care is required to estimate the shape parameter ξ within the global-local-tail shrinkage framework and we regard this as one of the salient contributions of the paper. For the GLT prior, we proposed an algorithm which combined the elliptical slice sampler [32] and the Hill estimator [23] from the extreme value theory which obviates the need for tuning any hyper-parameters. This automatic-tuning leads to learning the shape parameter ξ adaptive to the unknown sparsity level. Refer to §5 in the supplementary materials for more details.

REFERENCES

- [1] BAI, R. & GHOSH, M. (2018). High-dimensional multivariate posterior consistency under global–local shrinkage priors. *Journal of Multivariate Analysis* **167**, 157–170.
- [2] BAI, R. & GHOSH, M. (2018). On the beta prime prior for scale parameters in high-dimensional bayesian regression models. *arXiv preprint arXiv:1807.06539* .
- [3] BHADRA, A., DATTA, J., POLSON, N. G., WILLARD, B. et al. (2017). The horseshoe+ estimator of ultra-sparse signals. *Bayesian Analysis* **12**, 1105–1131.
- [4] BHADRA, A., DATTA, J., POLSON, N. G. & WILLARD, B. T. (2017). Lasso meets horseshoe. *arXiv preprint arXiv:1706.10179* .
- [5] BHATTACHARYA, A., CHAKRABORTY, A. & MALLICK, B. K. (2016). Fast sampling with gaussian scale mixture priors in high-dimensional regression. *Biometrika* , asw042.
- [6] BHATTACHARYA, A., PATI, D., PILLAI, N. S. & DUNSON, D. B. (2015). Dirichlet–laplace priors for optimal shrinkage. *Journal of the American Statistical Association* **110**, 1479–1490.
- [7] BÜHLMANN, P. & VAN DE GEER, S. (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media.
- [8] CARVALHO, C. M., POLSON, N. G. & SCOTT, J. G. (2009). Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*.
- [9] CARVALHO, C. M., POLSON, N. G. & SCOTT, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika* **97**, 465–480.
- [10] CASTILLO, I., SCHMIDT-HIEBER, J., VAN DER VAART, A. et al. (2015). Bayesian linear regression with sparse priors. *The Annals of Statistics* **43**, 1986–2018.
- [11] CASTILLO, I., VAN DER VAART, A. et al. (2012). Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics* **40**, 2069–2101.

- [12] CHICCOLI, C., LORENZUTTA, S. & MAINO, G. (1992). Concerning some integrals of the generalized exponential-integral function. *Computers & Mathematics with Applications* **23**, 13–21.
- [13] COLES, S., BAWA, J., TRENNER, L. & DORAZIO, P. (2001). *An introduction to statistical modeling of extreme values*, vol. 208. Springer.
- [14] EFRON, B. (2008). Microarrays, empirical bayes and the two-groups model. *Statistical science* , 1–22.
- [15] EFRON, B. (2010). The future of indirect evidence. *Statistical science: a review journal of the Institute of Mathematical Statistics* **25**, 145.
- [16] EFRON, B. (2012). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, vol. 1. Cambridge University Press.
- [17] EMBRECHTS, P., KLÜPPELBERG, C. & MIKOSCH, T. (2013). *Modelling extremal events: for insurance and finance*, vol. 33. Springer Science & Business Media.
- [18] FRIEDMAN, J., HASTIE, T. & TIBSHIRANI, R. (2001). *The elements of statistical learning*, vol. 1. Springer series in statistics New York.
- [19] GEORGE, E. I. & MCCULLOCH, R. E. (1995). Stochastic search variable selection. *Markov chain Monte Carlo in practice* **68**, 203–214.
- [20] GIERZ, G., HOFMANN, K., KEIMEL, K., LAWSON, J., MISLOVE, M. & SCOTT, D. (2003). Encyclopedia of mathematics and its applications.
- [21] GRIFFIN, J. E., BROWN, P. J. et al. (2010). Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis* **5**, 171–188.
- [22] HASTIE, T., TIBSHIRANI, R. & WAINWRIGHT, M. (2015). *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC.

- [23] HILL, B. M. (1975). A simple general approach to inference about the tail of a distribution. *The annals of statistics* , 1163–1174.
- [24] JEFFREYS, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **186**, 453–461.
- [25] JOHNDROW, J. E., ORENSTEIN, P. & BHATTACHARYA, A. (2017). Bayes shrinkage at gwas scale: Convergence and approximation theory of a scalable mcmc algorithm for the horseshoe prior. *arXiv preprint arXiv:1705.00841* .
- [26] JOHNSON, V. E. & ROSSELL, D. (2010). On the use of non-local prior densities in bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**, 143–170.
- [27] KARAMATA, J. (1933). Sur un mode de croissance régulière. théorèmes fondamentaux. *Bulletin de la Société Mathématique de France* **61**, 55–62.
- [28] MARIC, V. (2000). *Regular variation and differential equations*, vol. 1726. Springer Science & Business Media.
- [29] MARTIN, R., MESS, R., WALKER, S. G. et al. (2017). Empirical bayes posterior concentration in sparse high-dimensional linear models. *Bernoulli* **23**, 1822–1847.
- [30] MILGRAM, M. (1985). The generalized integro-exponential function. *Mathematics of computation* **44**, 443–458.
- [31] MITCHELL, T. J. & BEAUCHAMP, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association* **83**, 1023–1032.
- [32] MURRAY, I., PRESCOTT ADAMS, R. & MACKAY, D. J. (2010). Elliptical slice sampling .
- [33] PARK, T. & CASELLA, G. (2008). The bayesian lasso. *Journal of the American Statistical Association* **103**, 681–686.

- [34] PATI, D., BHATTACHARYA, A., PILLAI, N. S., DUNSON, D. et al. (2014). Posterior contraction in sparse bayesian factor models for massive covariance matrices. *The Annals of Statistics* **42**, 1102–1130.
- [35] PICKANDS III, J. et al. (1975). Statistical inference using extreme order statistics. *the Annals of Statistics* **3**, 119–131.
- [36] PIIRONEN, J., VEHTARI, A. et al. (2017). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics* **11**, 5018–5051.
- [37] POLSON, N. G. & SCOTT, J. G. (2010). Shrink globally, act locally: Sparse bayesian regularization and prediction. *Bayesian statistics* **9**, 501–538.
- [38] ROULEAU, M., PATEL, A., HENDZEL, M. J., KAUFMANN, S. H. & POIRIER, G. G. (2010). Parp inhibition: Parp1 and beyond. *Nature reviews cancer* **10**, 293.
- [39] SINGH, D., FEBBO, P. G., ROSS, K., JACKSON, D. G., MANOLA, J., LADD, C., TAMAYO, P., RENSHAW, A. A., D’AMICO, A. V., RICHIE, J. P. et al. (2002). Gene expression correlates of clinical prostate cancer behavior. *Cancer cell* **1**, 203–209.
- [40] SONG, Q. & LIANG, F. (2017). Nearly optimal bayesian shrinkage for high dimensional regression. *arXiv preprint arXiv:1712.08964* .
- [41] STEPHENS, P. J., TARPEY, P. S., DAVIES, H., VAN LOO, P., GREENMAN, C., WEDGE, D. C., NIK-ZAINAL, S., MARTIN, S., VARELA, I., BIGNELL, G. R. et al. (2012). The landscape of cancer genes and mutational processes in breast cancer. *Nature* **486**, 400.
- [42] STRATTON, M. R., CAMPBELL, P. J. & FUTREAL, P. A. (2009). The cancer genome. *Nature* **458**, 719.
- [43] TIPPING, M. E. (2001). Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research* **1**, 211–244.

- [44] VAN DER PAS, S., KLEIJN, B., VAN DER VAART, A. et al. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics* **8**, 2585–2618.
- [45] VAN DER PAS, S., SALOMOND, J.-B., SCHMIDT-HIEBER, J. et al. (2016). Conditions for posterior contraction in the sparse normal means problem. *Electronic journal of statistics* **10**, 976–1000.
- [46] VAN DER PAS, S., SZABÓ, B. & VAN DER VAART, A. (2016). How many needles in the haystack? adaptive inference and uncertainty quantification for the horseshoe. *arXiv preprint arXiv:1607.01892* .
- [47] VAN DER PAS, S., SZABÓ, B., VAN DER VAART, A. et al. (2017). Adaptive posterior contraction rates for the horseshoe. *Electronic Journal of Statistics* **11**, 3196–3225.
- [48] VAN DER PAS, S., SZABÓ, B., VAN DER VAART, A. et al. (2017). Uncertainty quantification for the horseshoe (with discussion). *Bayesian Analysis* **12**, 1221–1274.
- [49] VAN DER PAS, S. L., KLEIJN, B. J., VAN DER VAART, A. W. et al. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics* **8**, 2585–2618.
- [50] WILLYARD, C. (2018). New human gene tally reignites debate. *Nature* **558**, 354.
- [51] YANG, Y., WAINWRIGHT, M. J., JORDAN, M. I. et al. (2016). On the computational complexity of high-dimensional bayesian variable selection. *The Annals of Statistics* **44**, 2497–2532.
- [52] YOO, W. W. (2017). Contributed discussion to uncertainty quantification for the horseshoe by st\’ephanie van der pas, botond szab\’o and aad van der vaart. *arXiv preprint arXiv:1710.05987* .

- [53] ZHANG, R. & GHOSH, M. (2019). Ultra high-dimensional multivariate posterior contraction rate under shrinkage priors. *arXiv preprint arXiv:1904.04417* .

Supplementary Material to

Continuous shrinkage prior revisited: a collapsing behavior and remedy

Se Yoon Lee, Debdeep Pati, and Bani K. Mallick

Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX, U.S.A.

July 7, 2020

1. COLLAPSING BEHAVIOR OF HORSESHOE

1.1 The Horseshoe applied to a breast cancer data

The breast cancer data we use in this article is composed of a response vector and a design matrix, $(y, X) \in \mathfrak{R}^n \times \mathfrak{R}^{n \times p}$, obtained from $n = 729$ breast cancer patients and $p = 3,250$ genes. The i -th response value $y_i \in \mathfrak{R}$, $i = 1, \dots, n$, is the log-transformed overall survival (OS) time of the i -th subject such that all responses $\{y_i\}_{i=1}^n$ were quality assessed, integrated and processed with the help from disease experts and TCGA Biospecimen Core Resource [17]. Following a guideline from [17], subjects who have moderately long OS are considered in our study. A detailed clinical information of the dataset can be found in [17]. The minimum, mean, and maximum of OS are 84 days, 1,000 days (2.7 years), and 8,605 days (23 years), respectively. X is a column-standardized design matrix such that the ij -th element x_{ij} represents an expression levels of the j -th gene obtained from the i -th subject.

National Cancer Institute (NCI) defines OS as the length of time from either the date of diagnosis or the start of treatment for a disease, such as cancer, that patients diagnosed with the disease are still alive. In a clinical trial, measuring the OS is one way to see how well a new treatment

works. Therefore, OS is an indirect evidence of measuring about how strong the immune system of the patients. The histogram of $\{y_i\}_{i=1}^n$ and its Q-Q plot are displayed on the Figure 1. The Q-Q plot shows small deviation of the responses $\{y_i\}_{i=1}^n$ from normality.




Figure 1: Histogram of 729 log of overall survival times (left panel) and its Q-Q plot (right panel).

After centering the response vector y to avoid introducing an intercept term, our goal is to estimate β from the sparse high-dimensional linear regression (3). The primary objective of this study is then to discover two categories of small number of genes that may enhance the immune system of patients (positive sign of β_j) and genes who may undermine the immune system of patients (negative sign of β_j). As before, we use the Horseshoe $\pi_{\text{HS}}(\beta)$ for β , and Jeffrey's prior for σ . Posterior computation is executed by using `horseshoe(y = y, X = X, method.tau = "halfCauchy", method.sigma = "Jeffreys", burn = 10000, nmc = 10000, thin = 100)` where $y = y$ and $X = X$.

To investigate the behavior of the Horseshoe as the number of genes used increases, we constructed four datasets, $\mathcal{B}_1 = (y, X[:, 1 : 500])$, $\mathcal{B}_2 = (y, X[:, 1 : 1000])$, $\mathcal{B}_3 = (y, X[:, 1 : 2000])$, and $\mathcal{B}_4 = (y, X[:, 1 : 3250] = X)$, so that they have the same response vector y but the number of genes used in the design matrix are different; \mathcal{B}_1 , \mathcal{B}_2 , \mathcal{B}_3 , and \mathcal{B}_4 use 500, 1,000, 2,000, and 3,250 number of genes, respectively. The dataset \mathcal{B}_4 is the full dataset, and \mathcal{B} stands for breast. Figure 2 displays the stacked histograms of the column-wise correlations obtained from the design




Figure 2: Stacked histogram of the column-wise correlations of the design matrices from the four breast cancer datasets, \mathcal{B}_l , $l = 1, 2, 3, 4$, restricted on $[-0.6, -0.4]$ (left panel) and $[0.8, 1]$ (right panel).

matrices from the four datasets. Left and right histograms are obtained by confining the correlations to intervals $[-0.6, -0.4]$ and $[0.8, 1]$, respectively. We note from Figure 2 that as the number of genes used increases, the genome-wise correlations get substantially intensified. We report the results of the posterior inference by displaying the gene ranking plot, where the coefficients in $\beta = (\beta_1, \dots, \beta_p)^\top \in \mathbb{R}^p$ are ranked based on the absolute values of the posterior mean $\{\widehat{\beta}_j\}_{j=1}^p$, ordered from largest to smallest. Figure 3 displays the top 50 genes obtained by using the Horseshoe for each dataset \mathcal{B}_l , $l = 1, 2, 3, 4$. Table 1 summarizes top 10 genes along with their names, and directions which have been taken from the signs of the posterior means. The results are reasonable for \mathcal{B}_1 and \mathcal{B}_2 , but collapses when applied to \mathcal{B}_3 and \mathcal{B}_4 . Based on the Table 1, it turns out that the genes NGEF and FAM138F are found to be the most significant for the datasets \mathcal{B}_1 and \mathcal{B}_2 , respectively, and both genes have negative effects on the response OS. Figure 3 can be used for uncertainty quantification associated with the coefficients.

1.2 The GLT prior applied to a breast cancer data

The GLT prior is applied to the same four breast cancer data \mathcal{B}_l , $l = 1, 2, 3, 4$, constructed in Subsection 1.1. Recall that the Horseshoe collapses when applied to \mathcal{B}_l , $l = 3, 4$: see the bottom panels in the Figure 3. The Figure 4 and the Table 2 show the top 50 gene ranking plots and top 10 interesting genes obtained by using the GLT prior when applied to the four breast cancer datasets.

Table 3 summarizes the top 13 interesting genes selected by the GLT prior when applied to the

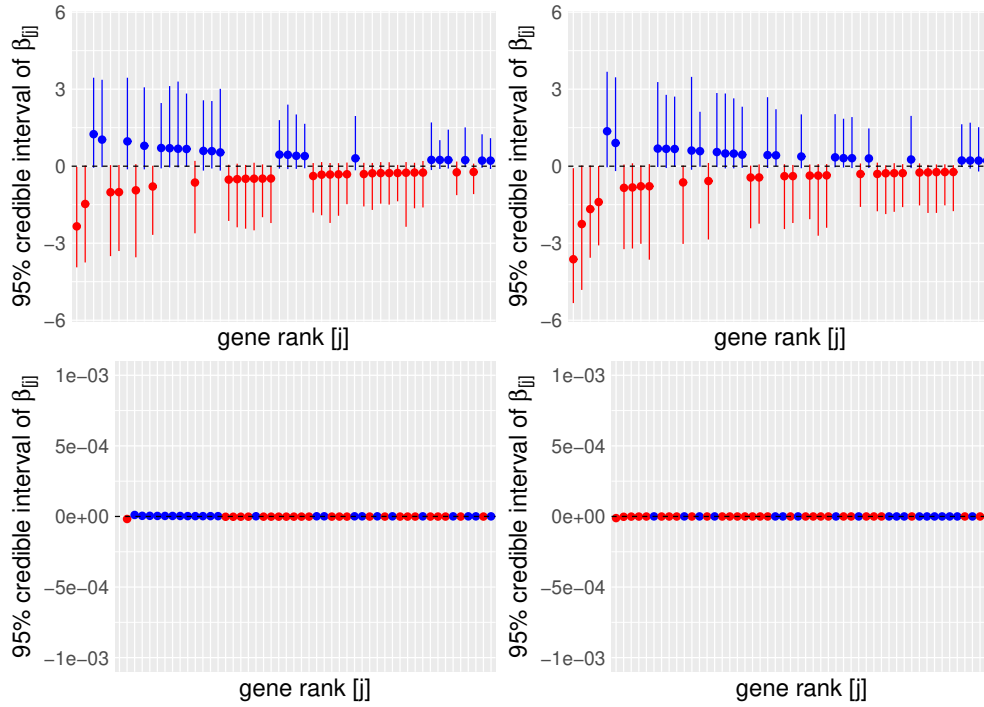


Figure 3: Top 50 genes obtained by the Horseshoe: \mathcal{B}_1 (top-left panel), \mathcal{B}_2 (top-right panel), \mathcal{B}_3 (bottom-left panel), and \mathcal{B}_4 (bottom-right panel). The dots \bullet and vertical bars represent the posterior means and 95% credible intervals, respectively. The colors blue and red represent plus and negative signs of posterior mean of β_j , respectively. Posterior means of τ corresponding to the four datasets are 0.10839 (\mathcal{B}_1), 0.06145 (\mathcal{B}_2), $3.65 \cdot 10^{-8}$ (\mathcal{B}_3), and $3.19 \cdot 10^{-9}$ (\mathcal{B}_4), respectively.

Table 1: Top 10 interesting genes selected by the Horseshoe when applied to \mathcal{B}_l , $l = 1, 2, 3, 4$

	1	2	3	4	5
\mathcal{B}_1	NGEF(-)	PLN(-)	C3orf59(+)	C21orf63(+)	LOC100130331(-)
\mathcal{B}_2	FAM138F(-)	SLC39A4(-)	PLN(-)	NGEF(-)	PCGF5(+)
\mathcal{B}_3	NA	NA	NA	NA	NA
\mathcal{B}_4	NA	NA	NA	NA	NA
	6	7	8	9	10
\mathcal{B}_1	FCGR2A(-)	HES4(+)	BCAP31(-)	GSTM1(+)	TOB2(-)
\mathcal{B}_2	HES4(+)	FCGR2A(-)	FCGR2C(-)	TOB2(-)	BCAP31(-)
\mathcal{B}_3	NA	NA	NA	NA	NA
\mathcal{B}_4	NA	NA	NA	NA	NA

NOTE: Contents of table is (gene name, direction). Genes with positive sign (+) may enhance the immune system of patients: however, genes with minus (-) may damage the immune system of patients. When the Horseshoe is applied to the datasets \mathcal{B}_3 and \mathcal{B}_4 , genes are unranked because the horseshoe estimator collapsed.

Table 2: Top 10 interesting genes selected by the GLT prior when applied to \mathcal{B}_l , $l = 1, 2, 3, 4$

	1	2	3	4	5
\mathcal{B}_1	NGEF(-)	C21orf63(+)	PLN(-)	C3orf59(+)	FCGR2A(-)
\mathcal{B}_2	FAM138F(-)	SLC39A4(-)	NGEF(-)	PCGF5(+)	PLN(-)
\mathcal{B}_3	FAM138F(-)	SLC39A4(-)	NGEF(-)	PLN(-)	COL7A1(-)
\mathcal{B}_4	FAM138F(-)	NSUN4(-)	COL7A1(-)	LOC150776(+)	NGEF(-)
	6	7	8	9	10
\mathcal{B}_1	BCAP31(-)	GSTM1(+)	LOC100130331(-)	TOB2(-)	ABCA17P(+)
\mathcal{B}_2	FCGR2A(-)	CRHR1(+)	TOB2(-)	GSTM1(+)	LOC150776(+)
\mathcal{B}_3	CRHR1(+)	FCGR2A(-)	RPLP1(+)	HES4(+)	TOB2(-)
\mathcal{B}_4	SMCHD1(+)	RPLP1(+)	HES4(+)	SLC37A2(-)	SLC39A4(-)

full breast cancer dataset \mathcal{B}_4 , and some references from the literature on oncology and genetics. The GLT prior discovered LOC150776 that has been less studied in the literature. As the direction of LOC150776 is positive (+), an over expression of LOC150776 may enhance the immune system of breast cancer patients. Interestingly, the GLT prior identified the famous superman gene BHLHE41: it is known that the genetic variant of BHLHE41 provides a greater resistance to the effects of sleep deprivation, possibly enhancing the immune system [24].

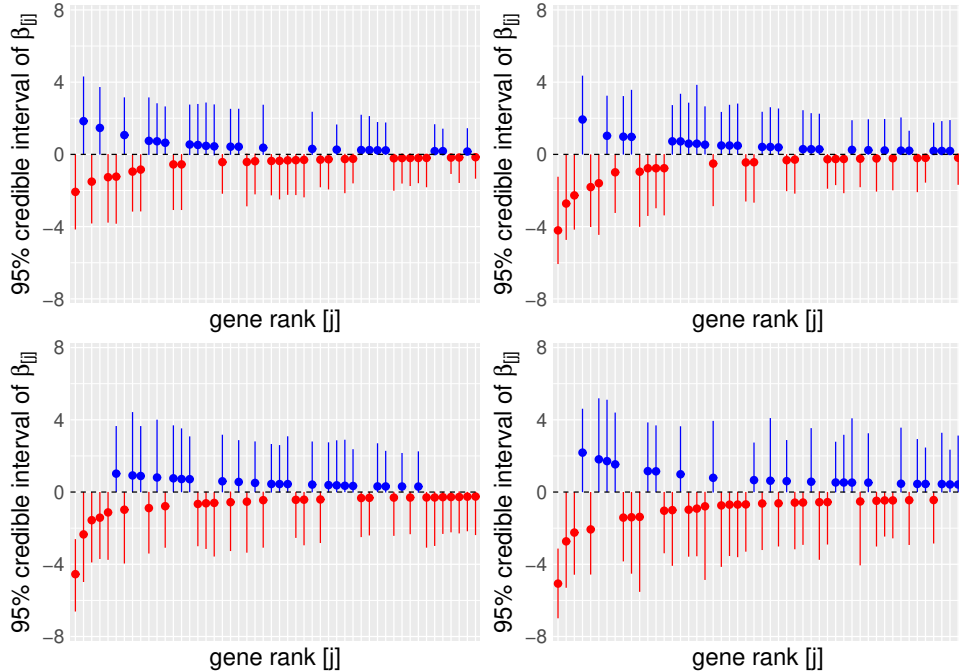


Figure 4: Top 50 gene ranking plots obtained by the GLT prior: \mathcal{B}_1 (top-left panel), \mathcal{B}_2 (top-right panel), \mathcal{B}_3 (bottom-left panel), and \mathcal{B}_4 (bottom-right panel). Posterior means of (τ, ξ) corresponding to the four datasets are $(0.00436, 2.188)$ (\mathcal{B}_1), $(0.00221, 2.230)$ (\mathcal{B}_2), $(0.00135, 2.382)$ (\mathcal{B}_3), and $(0.00135, 2.922)$ (\mathcal{B}_4), respectively.

2. CURVE FITTING STUDY

2.1 Simulated curves

Consider two functions f on domain \mathcal{D} from what data is generated: the sinc curve $f(x) = \text{sinc}(x) = (\sin x)/x$ on $\mathcal{D} = (-20, 20)$, and a flat curve $f(x) = (5x - 3)^3 \cdot \mathcal{I}(x > 3/5)$ on $\mathcal{D} = (0, 1)$. We uniformly sampled n -inputs $\{x_i\}_{i=1}^n$ from domain \mathcal{D} , and let $y_i = f(x_i) + \sigma_0 \epsilon_i$, $\epsilon_i \sim \mathcal{N}_1(0, 1)$, $i = 1, \dots, n$, with $\sigma_0 = 0.15$, to generate n -pair $\{(y_i, x_i)\}_{i=1}^n$.

To estimate f given the n -pair $\{(y_i, x_i)\}_{i=1}^n$ the sparse Gaussian kernel regression [4, 28] is used as: $y_i = f_n(x_i) + \sigma \epsilon_i$, $\epsilon_i \sim \mathcal{N}_1(0, 1)$, $i = 1, \dots, n$, such that $f_n(\cdot) = \alpha + \sum_{j=1}^n \beta_j \mathcal{K}(\cdot, x_j)$, $\alpha \in \mathfrak{R}$, and $\beta = (\beta_1, \dots, \beta_n)^\top \in \mathfrak{R}^n$ is sparse, with the Gaussian kernel for \mathcal{K} [4]. We use the flat prior for α [18], and the Jeffrey's prior for σ . Sparsity on β is imposed by the Horseshoe or the GLT prior. For each test curve, we generated $n = 100$ pairs, and report the median of all average mean squared error (AMSE) [32] across 100 replicated pairs $\{(y_i, x_i)\}_{i=1}^{n=100}$. AMSE is defined by $\sum_{i=1}^n \{\hat{f}_n(x_i) - f(x_i)\}^2 / n$, where $\hat{f}_n(x) = E[\alpha + \sum_{j=1}^n \beta_j \mathcal{K}(x, x_j) | y]$ is a posterior mean of

Table 3: Top 13 interesting genes selected by the GLT prior when applied to \mathcal{B}_4

Rank	Gene (direction)	Note	References
1	FAM138F(-)	Increasing a risk of breast and ovarian cancer	[11, 27]
2	NSUN4(-)	Related with ovarian and prostate cancer	[15]
3	COL7A1(-)	Related with cell migration (metastasis)	[33]
4	LOC150776(+)	Less studied in oncology and genetics	
5	NGEF(-)	Related with obesity-related diseases	[31]
6	SMCHD1(+)	Important in regulation	[14]
7	RPLP1(+)	Important in protein synthesis	[9]
8	HES4(+)	Gene knockdown increases a brain disease	[1]
9	SLC37A2(-)	Negatively related with survival probability	
10	SLC39A4(-)	Negatively related with survival probability	[14]
11	MFRP(-)	Related with ovarian cancer	
12	ARSA(+)	Positively related with survival probability	
13	BHLHE41(+)	High recovery from fatigue or short sleep	[24]

$f_n(x)$ at x .

For the sinc test curve, the median AMSE obtained by the Horseshoe and the GLT priors are 0.00393 and 0.00385, respectively. For the flat test curve, the median of AMSE obtained by using the Horseshoe and the GLT prior are 0.00490 and 0.00382, respectively. See Figure 5 for one of the 100 replicates.

2.2 Real curves

The sparse Gaussian kernel regression is applied to four example curves: circadian rhythm curve of gene expression of PER2 from colon tissue, light-curve from an eclipsing binary star system, fossil data, and LIDAR data. The number of observations for the four data are 100, 377, 106, and 221, respectively. The circadian rhythm data and light-curve data can be obtained from the website <http://circadb.hogeneschlab.org> and <https://www.eso.org>, respectively. The fossil data and the LIDAR data can be downloaded from R package `SemiPar`. See Figure 6 for the results: the results are virtually indistinguishable.

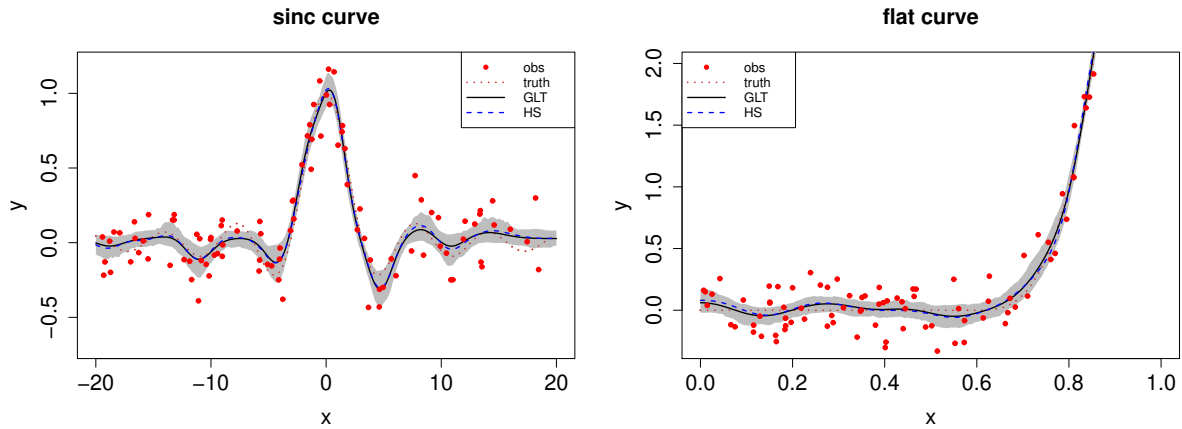


Figure 5: Two simulated curves fitted by the sparse Gaussian kernel regression. The red dot and red dotted curve represent observation and the truth f . The black curve and blue dotted curve represent the posterior mean of $f_n(x)$ at x obtained by using the GLT prior and the Horseshoe, respectively. The shaded region depicts the pointwise 95% credible interval obtained by using the GLT prior.

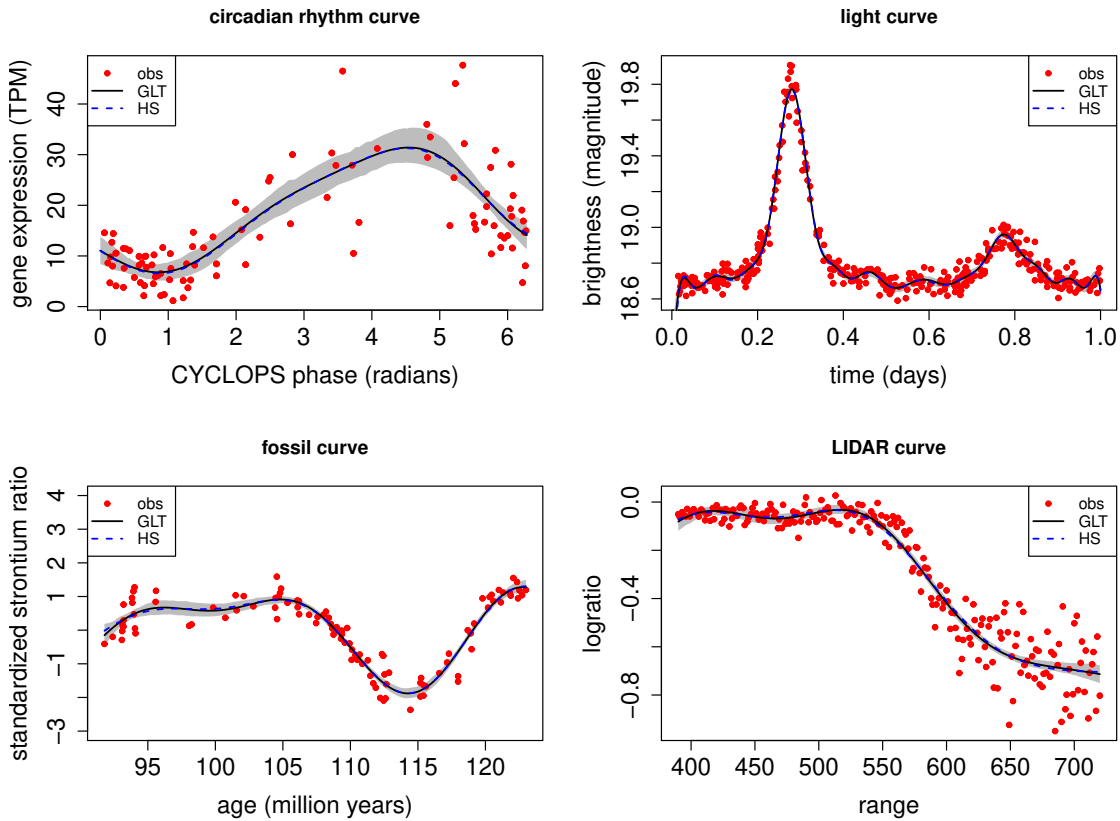


Figure 6: Four real curves fitted by the sparse Gaussian kernel regression.

3. ADDITIONAL PROPERTIES OF THE HORSESHOE AND GLT PRIOR

3.1 Marginal density of Horseshoe

Lemma 1 (Marginal density of Horseshoe).

(a) Assume $\beta|\lambda, \tau \sim \mathcal{N}_1(0, \tau^2\lambda^2)$, $\lambda \sim \mathcal{C}^+(0, 1)$, and $\tau > 0$. Then:

$$\pi_{HS}(\beta|\tau) = \int \mathcal{N}_1(\beta|0, \tau^2\lambda^2)\pi(\lambda)d\lambda = K_{HS}e^{Z_{HS}(\beta)}E_1\{Z_{HS}(\beta)\}, \quad (1)$$

where $K_{HS} = 1/(\tau 2^{1/2}\pi^{3/2})$ and $Z_{HS}(\beta) = \beta^2/(2\tau^2)$. $E_1(x) = \int_1^\infty e^{-xt}t^{-1}dt$, $x \in \mathfrak{R}$, is the exponential integral function.

(b) Assume $\lambda \sim \mathcal{C}^+(0, 1)$, $\kappa = 1/(1 + \tau^2\lambda^2) \in (0, 1)$, and $\tau > 0$. Then:

$$\pi_{HS}(\kappa|\tau) = \frac{\tau}{\pi} \cdot \frac{\kappa^{-1/2}(1 - \kappa)^{-1/2}}{1 - (1 - \tau^2)\kappa}. \quad (2)$$

3.2 The GLT prior for high-dimensional linear regression

Consider a high-dimensional linear regression

$$y = X\beta + \sigma\epsilon, \quad \epsilon \sim \mathcal{N}_n(0, I_n) \quad \text{and} \quad \beta \text{ is sparse}, \quad (3)$$

where X is a n -by- p design matrix ($n \ll p$ or $n \approx p$ with both n and p are large). The goal is to recover the p -dimensional coefficients vector $\beta = (\beta_1, \dots, \beta_p)^\top$ in a fully Bayesian way.

In the main paper, we proposed the GLT prior for the parameter β , $\pi_{GLT}(\beta)$, whose hierarchical formulation is given by:

$$\beta_j|\lambda_j, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2\sigma^2), \quad \sigma^2 \sim \pi(\sigma^2) \propto 1/\sigma^2, \quad (j = 1, \dots, p), \quad (4)$$

$$\lambda_j|\tau, \xi \sim \mathcal{GPD}(\tau, \xi), \quad (j = 1, \dots, p), \quad (5)$$

$$\tau|\xi \sim \mathcal{IG}(p/\xi + 1, 1), \quad (6)$$

$$\xi \sim \log \mathcal{N}(\mu, \rho^2)\mathcal{I}_{(1/2, \infty)}, \quad \mu \in \mathfrak{R}, \rho^2 > 0. \quad (7)$$

Note that μ and ρ^2 are hyper-parameters which typically require an expert-tuning. In this Supplementary material, we provide a full description of a posterior computation using the GLT prior under the high-dimensional linear regression. Eventually, the proposed posterior computation does NOT require any tuning procedure for the hyper-parameters, and the shape parameter ξ is adaptively learned based on the sparsity level associated with the given data (y, X) .

3.3 Learnability of the shape parameter ξ using the GLT prior

Since the shape parameter ξ is the furthest from the data y within the hierarchy, it is important to check properness of the posterior distribution of ξ , $\pi(\xi|y) \propto f(y|\xi) \cdot \pi(\xi)$. For that, one should prove that the evidence (marginal likelihood) $m(y) = \int f(y|\xi) \cdot \pi(\xi) d\xi = \int \int \int \int \int \pi(y, \beta, \sigma^2, \lambda, \tau, \xi) d\beta d\sigma^2 d\lambda d\tau d\xi$ is finite for all values $y \in \mathbb{R}^n$, which is not trivial. Instead, we demonstrate properness of two posterior densities: (a) the full conditional posterior distribution $\pi(\xi|y) = \pi(\xi|\lambda, \tau)$ from the hierarchy of the GLT prior, and (b) the posterior distribution $\pi(\xi|y, \beta, \tau, \lambda)$ under a univariate hierarchy without covariates:

Lemma 2. (*Learnability of ξ using the GLT prior*)

(a) Assume $\beta \sim \pi_{GLT}(\beta)$ (4) – (7). Then proportional part of full conditional posterior for ξ is represented as:

$$\pi(\xi|y) \propto \mathcal{V}_p(\xi) \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi), \quad \mathcal{V}_p(\xi) = \frac{\pi^{p/2}}{\Gamma(p/\xi + 1)} \prod_{j=1}^p r_j(\xi), \quad (8)$$

where $\{r_j(\xi)\}_{j=1}^p = (\tau + \xi\lambda_j)^{-(1/\xi+1)}$ and $\pi(\xi|y)$ is proper on $(1/2, \infty)$. Here, \mathcal{V} stands for volume.

(b) Assume $y|\beta \sim \mathcal{N}_1(\beta, 1)$, $\beta|\lambda \sim \mathcal{N}_1(0, \lambda^2)$, $\lambda|\tau, \xi \sim \mathcal{GPD}(\tau, \xi)$, $\tau|\xi \sim \mathcal{IG}(1/\xi + 1, 1)$. Let $\pi(\xi)$ be any proper density of ξ supported on $(1/2, \infty)$, i.e., $\int_{1/2}^{\infty} \pi(\xi) d\xi = 1$. Then $\pi(\xi|y, \beta, \tau, \lambda)$ is proper on $(1/2, \infty)$.

Interestingly, the likelihood part of the full conditional posterior density $\pi(\xi|y)$ (8) has a nice

geometric interpretation: if $\xi = 2$ then the value of $\mathcal{V}_p(2)$ of the density is the volume of a p -dimensional ellipsoid with p -radii $\{r_j(2) = (\tau + 2\lambda_j)^{-(3/2)}\}_{j=1}^p$.

4. SIMULATION STUDIES WITH VARIANTS OF THE HORSESHOE

We conducted replicated study under the high-dimensional regression (3) with $n = 100$ and $p = 500$. Simulation environments are coincided with the three scenarios described in the Section 6 in the main paper. (Each of the scenario 1, 2, and 3 vary sparsity level, correlation ρ associated with design matrix, and signal-to-noise (SNR) ratio, while other simulation setting fixed.) Here, we additionally investigate three variant versions [3, 25, 29] of the Horseshoe along with the main target Horseshoe [5]:

Truncated horseshoe [29].

$$\beta_j | \lambda_j, \tau, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2 \tau^2 \sigma^2), \lambda_j \sim \mathcal{C}^+(0, 1), \tau \sim \mathcal{TC}^+(0, 1)_{(1/p, \infty)}, \quad (j = 1, \dots, p).$$

The $\mathcal{TC}^+(0, 1)_{(1/p, \infty)}$ is the unit-scaled half-Cauchy distribution truncated from below by $1/p$. The R function `horseshoe` within the R package `horseshoe` provides an option to use this setting by specifying `method.tau = "truncatedCauchy"`.

Horseshoe-plus [3].

$$\beta_j | \lambda_j, \sigma^2 \sim \mathcal{N}_1(0, \lambda_j^2 \sigma^2), \lambda_j | \eta_j, \tau \sim \mathcal{C}^+(0, \eta_j \tau), \eta_j, \tau \sim \mathcal{C}^+(0, 1), \quad (j = 1, \dots, p).$$

Note that the Horseshoe-plus is characterized by a further half-Cauchy mixing variable η_j embedded to the local-scales λ_j .

Regularized horseshoe [25].

$$\begin{aligned} \beta_j | \tilde{\lambda}_j, \tau \sim \mathcal{N}_1(0, \tau^2 \tilde{\lambda}_j^2), \quad \tilde{\lambda}_j^2 &= \frac{c^2 \lambda_j^2}{c^2 + \tau^2 \lambda_j^2}, \quad (j = 1, \dots, p), \\ \lambda_j, \tau \sim \mathcal{C}^+(0, 1), \quad c^2 &\sim \mathcal{IG}(\nu/2, \nu s^2/2), \quad (j = 1, \dots, p), \end{aligned}$$

where $\nu, s^2 > 0$ are hyper-parameters: we shall simply fix them to be 1.

For evaluation criteria, given that the truth $\beta_0 = (\beta_{0,1}, \dots, \beta_{0,q}, \beta_{0,q+1}, \dots, \beta_{0,p})^\top$ is specified by $\beta_{0,1} = \dots = \beta_{0,q} = 1$ (q unit signals) and $\beta_{0,q+1} = \dots = \beta_{0,p} = 0$ ($p - q$ noises), we separately report median of the following quantities obtained by 50 replicates:

$$\text{MSE} = \frac{1}{p} \sum_{j=1}^p (\hat{\beta}_j - \beta_{0,j})^2, \quad \text{MSE}_S = \frac{1}{q} \sum_{j=1}^q (\hat{\beta}_j - 1)^2, \quad \text{and} \quad \text{MSE}_N = \frac{1}{p - q} \sum_{j=q+1}^p (\hat{\beta}_j)^2.$$

The MSE measures overall accuracy of estimation for the coefficients induced by a prior, which can be dissected by two components: (1) MSE for signal part (MSE_S) measuring signal recovery ability and (2) MSE for noise parts (MSE_N) measuring noise shrinking ability. We emphasize that when collapse takes place (that is, the posterior means $\hat{\beta}_j$ are nearly zeros), then MSE_S and MSE_N will be close to 1 and 0, respectively. In this circumstance, the total MSE is not a reasonable evaluation criteria.

Figure 7 displays the simulation results: scenario 1 (top three panels); scenario 2 (middle three panels); and scenario 3 (bottom three panels). The followings are summaries based on the results:

1. Under the scenario 1, we see that the truncated horseshoe prior [30] suffers from the similar collapse observed in the Horseshoe [5] when sparsity level is larger than certain threshold. The total does not bring out this phenomenon.
2. Under the scenario 1 with ultra sparsity regime (where the sparsity level q/p is between 0.002 and 0.024), all considered prior performs reasonably well, while the signal recovery ability of the GLT prior is marginally getting better as the sparsity level increases.
3. Under the scenario 1 with moderate sparsity regime (where the sparsity level q/p is between 0.034 and 0.1), (i) the regularized horseshoe [25] outperforms others in terms of MSE_S , while (ii) the GLT prior outperforms others in terms of MSE_N .
4. Under the scenario 2, the Horseshoe [5] and the truncated horseshoe [30] outperform other priors in terms of MSE, while the GLT prior outperforms in terms of MSE_S .

5. Under the scenario 3, the Horseshoe [5] and the truncated horseshoe [30] outperform other priors in terms of MSE, while the GLT prior outperforms in terms of MSE_S when SNR is 2.

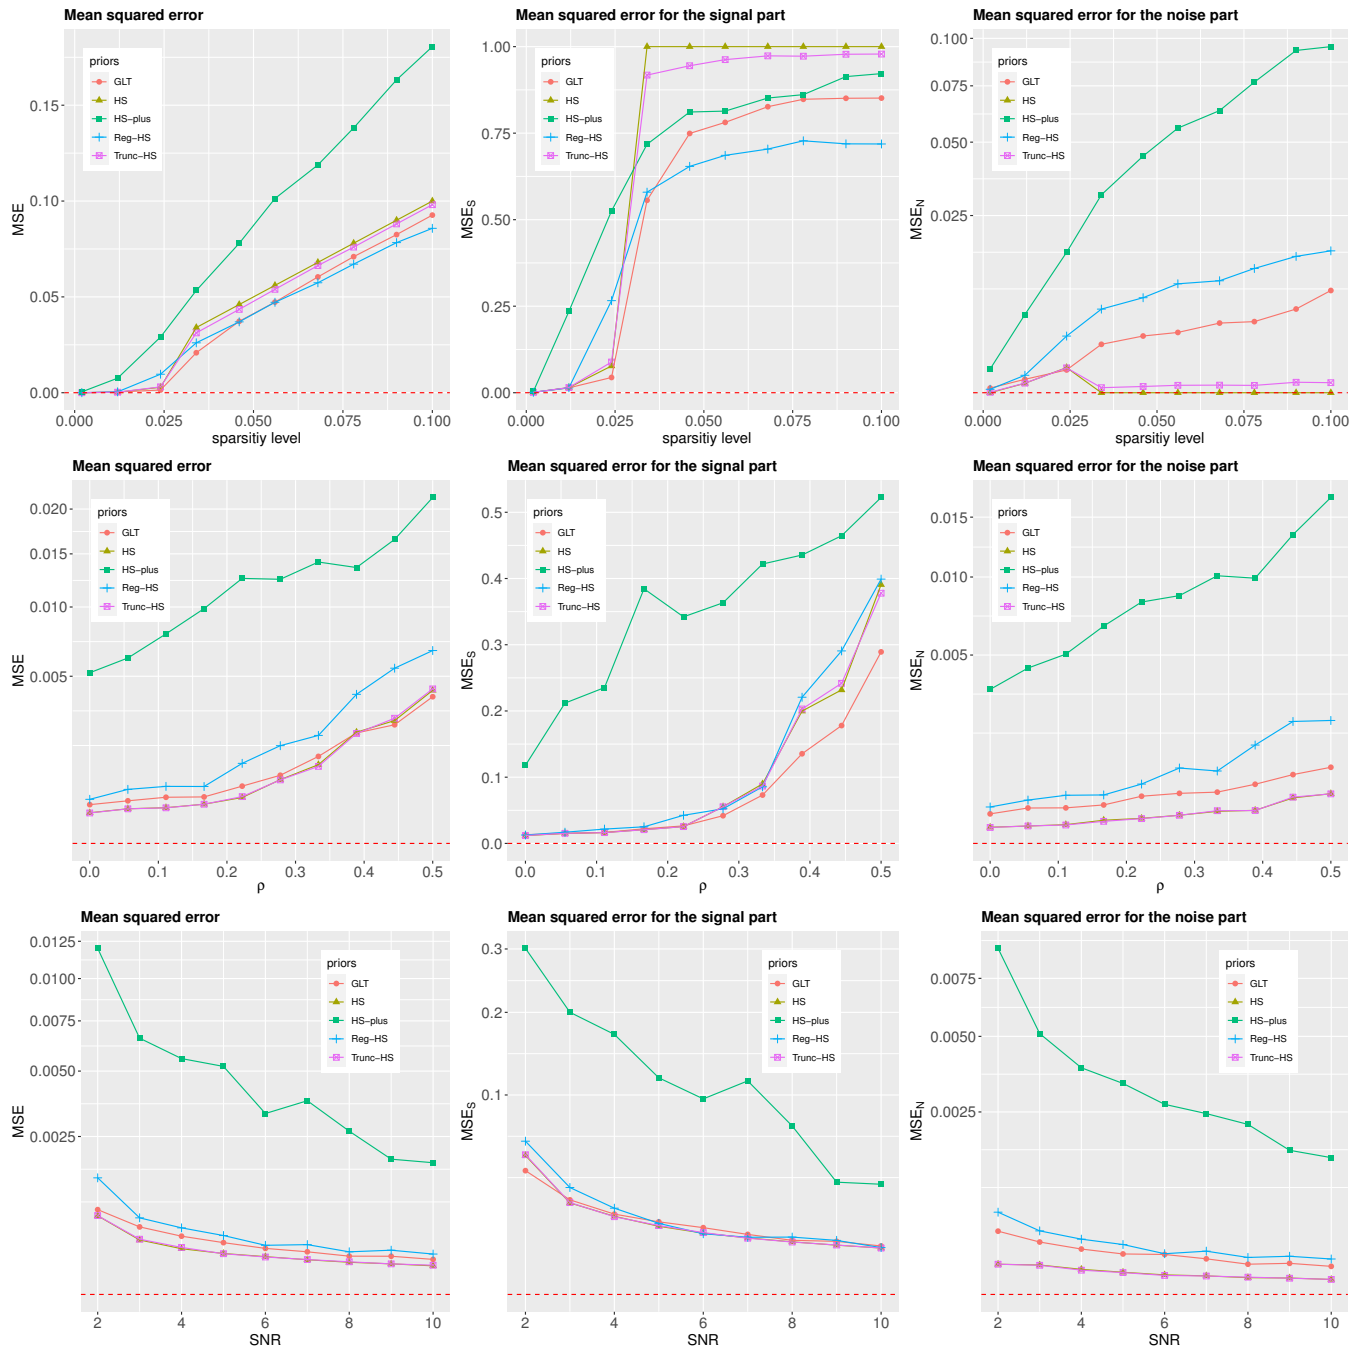


Figure 7: Simulation results under the three scenarios: scenario 1 (top panels); scenario 2 (middle panels); and scenario 3 (bottom panels). Metrics measured are MSE (left panels), MSE_S (center panels), and MSE_N (right panels).

5. POSTERIOR COMPUTATION

In this section, we details several aspects of the posterior computation using GLT prior.

5.1 Gibbs sampler

Consider $y = X\beta + \sigma^2\epsilon$, $\epsilon \sim \mathcal{N}_n(0, I_n)$ (3), $\sigma^2 \sim \pi(\sigma^2) \propto 1/\sigma^2$, and $\beta \sim \pi_{\text{GLT}}(\beta)$ (4) – (7).

The posterior distribution of unknown parameters $\Omega = (\beta, \sigma^2, \lambda, \tau, \xi) \in \mathfrak{R}^p \times (0, \infty) \times (0, \infty)^p \times (0, \infty) \times (1/2, \infty)$ is obtained in the following. The full joint posterior distribution, $\pi(\Omega|y)$, is proportional to

$$\begin{aligned} & \mathcal{N}_n(y|X\beta, \sigma^2 I_n) \mathcal{N}_p(\beta|0, \sigma^2 \Lambda) \pi(\sigma^2) \left\{ \prod_{j=1}^p \pi(\lambda_j|\tau, \xi) \right\} \pi(\tau, \xi), \quad \Lambda = \text{diag}(\lambda_1^2, \dots, \lambda_p^2) \in \mathfrak{R}^{p \times p} \\ & \propto \mathcal{N}_n(y|X\beta, \sigma^2 I_n) \mathcal{N}_p(\beta|0, \sigma^2 \Lambda) \pi(\sigma^2) \left\{ \prod_{j=1}^p \mathcal{GPD}(\lambda_j|\tau, \xi) \right\} \mathcal{IG}(\tau|p/\xi + 1, 1) \log \mathcal{N}(\xi|\mu, \rho^2) \mathcal{I}_{(1/2, \infty)}(\xi). \end{aligned}$$

Since full joint posterior distribution $\pi(\Omega|y)$ is not in a closed form, we develop a Markov chain Monte Carlo to simulate Ω from this full joint posterior distribution. The following algorithm provides a Gibbs sampler which utilizes the conditional independence structure in the hierarchical formation.

Step 1. Sample β from conditional posterior

$$\pi(\beta|-) \sim \mathcal{N}_p(\Sigma X^T y, \sigma^2 \Sigma), \quad \Sigma = (X^T X + \Lambda^{-1})^{-1} \in \mathfrak{R}^{p \times p}.$$

Step 2. Sample σ^2 from conditional posterior

$$\pi(\sigma^2|-) \sim \mathcal{IG}\left(\frac{n+p}{2}, \frac{\|y - X\beta\|_2^2 + \beta^T \Lambda^{-1} \beta}{2}\right).$$

Step 3. Update λ_j , $j = 1, \dots, p$, independently using slice sampler [22] within the Gibbs sampler. Proportional part of full conditional posterior is

$$\pi(\lambda_j|-) \propto \frac{1}{\lambda_j} \exp\left(-\frac{\beta_j^2}{2\sigma^2 \lambda_j^2}\right) \cdot \left(1 + \frac{\xi \lambda_j}{\tau}\right)^{-(1/\xi+1)}. \quad (9)$$

Step 4. Update τ using slice sampler [22] within the Gibbs sampler. Proportional part of full conditional posterior is

$$\pi(\tau|-) \propto \tau^{-2} \exp(-1/\tau) \cdot \prod_{j=1}^p (\tau + \xi \lambda_j)^{-(1/\xi+1)}. \quad (10)$$

Step 5. Update ξ using elliptical slice sampler [21] after variable change $\eta = \log \xi$ within the Gibbs sampler. Proportional part of full conditional posterior is

$$\pi(\xi|-) \propto \mathcal{V}_p(\xi) \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi), \quad (11)$$

where $\mathcal{V}_p(\xi) = \{\Gamma(p/\xi+1)\}^{-1} \pi^{p/2} \prod_{j=1}^p r_j(\xi)$ with $r_j(\xi) = (\tau + \xi \lambda_j)^{-(1/\xi+1)}$, $j = 1, \dots, p$.

5.2 Slice sampler implementation in *Step 3* and *Step 4*

Slice sampler [22] is a popular technique to adapt the step-size of a MCMC algorithm and is based on the local property of the target density. The basic idea is parameter expansion which involves intentional introduction of auxiliary variables [7]. Finding an appropriate parameter expansion depends on the functional form of the target density.

Let $j \in \{1, \dots, p\}$. To implement the slice sampler in the *Step 3* (9), first use change of variable, $\gamma_j = \lambda_j^2$, to get

$$\begin{aligned} \pi(\gamma_j|-) &\propto \gamma_j^{-1} \exp(-m_j/\gamma_j) \cdot (\tau + \xi \sqrt{\gamma_j})^{-(1/\xi+1)} \\ &= \gamma_j^{-1} \exp(-m_j/\gamma_j) \cdot (\sqrt{\gamma_j})^{-(1/\xi+1)} (\sqrt{\gamma_j})^{(1/\xi+1)} \cdot (\tau + \xi \sqrt{\gamma_j})^{-(1/\xi+1)} \\ &= \gamma_j^{-(1/\xi+1)/2-1} \exp(-m_j/\gamma_j) \cdot (\xi + \tau \cdot \gamma_j^{-1/2})^{-(1/\xi+1)} \\ &\propto \mathcal{IG}\{\gamma_j|(1/\xi+1)/2, m_j\} \cdot g(\gamma_j), \end{aligned} \quad (12)$$

where $m_j = \beta_j^2/(2\sigma^2)$ and $g(\gamma_j) = (\xi + \tau \cdot \gamma_j^{-1/2})^{-(1/\xi+1)}$. Note that the function $u_j = g(\gamma_j)$ is increasing on $(0, \infty)$, and its inverse function is $\gamma_j = g^{-1}(u_j) = [\tau/\{u_j^{-(\xi/(1+\xi))} - \xi\}]^2$. Now, consider a density, $\pi(\gamma_j, u_j|-) \propto \mathcal{IG}\{\gamma_j|(1/\xi+1)/2, m_j\} \cdot \mathcal{I}_{(0, g(\gamma_j))}(u_j)$. Then we can show that

$\int \pi(\gamma_j, u_j | -) du_j = \pi(\gamma_j | -)$, which means that $\pi(\gamma_j, u_j | -)$ is a valid parameter expansion of (12). Actual sampling is executed on $\pi(\gamma_j, u_j | -)$ using the Gibbs sampler: (i) $u_j | \gamma_j, - \sim \pi(u_j | \gamma_j, -) = \mathcal{U}(0, g(\gamma_j))$ and (ii) $\gamma_j | u_j, - \sim \pi(\gamma_j | u_j, -) = \mathcal{IG}\{\gamma_j | (1/\xi + 1)/2, m_j\} \cdot \mathcal{I}_{(g^{-1}(u_j), \infty)}(\gamma_j)$. After the Gibbs sampling, transform back to $\lambda_j = \sqrt{\gamma_j}$.

To implement the slice sampler in the *Step 4*, note from (10):

$$\pi(\tau | -) \propto \mathcal{IG}(\tau | 1, 1) \cdot \prod_{j=1}^p g_j(\tau), \quad (13)$$

where $g_j(\tau) = (\tau + \xi\lambda_j)^{-(1/\xi+1)}$, $j = 1, \dots, p$. Note that p -functions $v_j = g_j(\tau)$, $j = 1, \dots, p$, are decreasing on $(0, \infty)$, and their inverse functions are $\tau = g_j^{-1}(v_j) = v_j^{-(\xi/(1+\xi))} - \xi\lambda_j$, $j = 1, \dots, p$. Now, consider a density: $\pi(\tau, v_1, \dots, v_p | -) \propto \mathcal{IG}(\tau | 1, 1) \cdot \prod_{j=1}^p \mathcal{I}_{(0, g_j(\tau))}(v_j)$. Then we have $\int \dots \int \pi(\tau, v_1, \dots, v_p | -) dv_1 \dots dv_p = \pi(\tau | -)$ and hence $\pi(\tau, v_1, \dots, v_p | -)$ is a valid parameter expansion of (13). Actual sampling is executed on $\pi(\tau, v_1, \dots, v_p | -)$ using the Gibbs sampler:

$$\begin{aligned} v_j | \tau, v_{-j}, - &\sim \pi(v_j | \tau, v_{-j}, -) = \mathcal{U}(0, g_j(\tau)), \quad (j = 1, \dots, p), \\ \tau | v_1, \dots, v_p, - &\sim \mathcal{IG}(\tau | 1, 1) \cdot \mathcal{I}_{(0, \min\{g_1^{-1}(v_1), \dots, g_p^{-1}(v_p)\})}(\tau), \end{aligned} \quad (14)$$

where in (14), v_{-j} represents the collection of $\{v_j\}_{j=1}^p$ except for v_j . Note also that each full conditional posterior distribution $\pi(v_j | \tau, v_{-j}, -)$, $j = 1, \dots, p$, does not depend on v_{-j} , i.e., $\pi(v_j | \tau, v_{-j}, -) = \pi(v_j | \tau, -)$ and hence it is possible to parallelize the update of $\{v_j\}_{j=1}^p$.

5.3 Summary of the Hill estimator

We briefly explain the Hill estimator which plays a central role in hyper-parameter specification of μ . For notational coherence, we use the Greek letter λ to describe a random quantity. Suppose $\lambda = (\lambda_1, \dots, \lambda_p)^T \in (0, \infty)^p$ is p -dimensional random variables from a strongly stationary process whose marginal distribution is F such that its tail distribution is regularly varying with the tail-index $1/\xi$ (hence, the corresponding shape parameter is ξ). To be specific, the tail distribution is

described as $\bar{F}(\lambda) = 1 - F(\lambda) = L(\lambda) \cdot \lambda^{-1/\xi}$ for some $\xi > 0$ where L is a slowly varying function [8, 26]. Denote its order statistics by $\lambda_{(1)} \geq \dots \geq \lambda_{(p)}$. Then the Hill estimator [12] is based on k upper order statistics:

$$\widehat{\xi}_k(\lambda) = \frac{1}{k-1} \sum_{j=1}^{k-1} \log \left(\frac{\lambda_{(j)}}{\lambda_{(k)}} \right), \quad \text{for } 2 \leq k \leq p. \quad (15)$$

It is known that the Hill estimator (15) is a consistent estimator for ξ , i.e., $\widehat{\xi}_k(\lambda) \rightarrow \xi$ in probability, if $p \rightarrow \infty$, $k \rightarrow \infty$, and $k/p \rightarrow 0$ [8, 10, 26]. Empirically it is known that the Hill estimator may work effectively when F is of Pareto type [8, 16]. (See Fig 1 in [8].)

Suppose we have p number of observations λ , possibly generated from the above distribution F . In practice, the Hill estimator is used as follows. First, calculate the estimator $\widehat{\xi}_k(\lambda)$ at each integer $k \in \{2, \dots, p\}$, and then plot the ordered pairs $\{(k, \widehat{\xi}_k(\lambda))\}_{k=2}^p$: the resulting plot is called the Hill plot (See the Figure 6.4.3 of [10]). Then, select from the set of Hill estimators $\{\widehat{\xi}_k(\lambda)\}_{k=2}^p$ which are stable (roughly constant) with respect to k : then, the stable value(s) are regarded as reasonable estimate(s) for the shape parameter ξ [8]. Typically, the Hill plot may display high variability when k is close to 2 or p . As a practical remedy, one may disregard the first or last few of the estimates: the values $\widehat{\xi}_k(\lambda)$ that are evaluated at integers $k \in \{k_L, \dots, k_U\}$, $2 < k_L < k_U < p$, are considered to be monitored where the integers k_L and k_U are decided by user.

5.4 Hyper-parameter specification of μ and ρ^2

Suppose we are at the *Step 5* of the s -th iteration of the Gibbs sampler described in Subsection 5.1. At this moment, we already acquired posterior realizations, $\lambda^{(s+1)} = (\lambda_1^{(s+1)}, \dots, \lambda_p^{(s+1)})^T$ and $\tau^{(s+1)}$, that are sampled from the previous steps, *Step 3* and *Step 4*, respectively. Treating the indicator $\mathcal{I}_{(1/2, \infty)}(\xi)$ in (11) as a part of likelihood, consider sampling $\xi^{(s+1)}$ from the density;

$$\xi^{(s+1)} \sim \pi(\xi | -) = \pi(\xi | \lambda^{(s+1)}, \tau^{(s+1)}) \propto \mathcal{L}(\xi) \cdot \log \mathcal{N}_1(\xi | \mu, \rho^2), \quad (16)$$

where $\mathcal{L}(\xi) = \mathcal{V}_p(\xi)\mathcal{I}_{(1/2,\infty)}(\xi)$. Henceforth, the basic idea is to strictly obey the philosophy of Gibbs sampler: as long as we are to sample $\xi^{(s+1)} \sim \pi(\xi| -)$, every latent variables except for ξ are treated as observed variables, including $\lambda^{(s+1)}$ and $\tau^{(s+1)}$.

To start with, we choose a small value of the hyper-parameter ρ^2 so that the prior $\pi(\xi) = \log \mathcal{N}_1(\xi|\mu, \rho^2)$ in (16) is highly concentrated around its prior mean $E[\xi] = \exp(\mu + \rho^2/2) \approx \exp(\mu)$. Hence, a future state $\xi^{(s+1)}$ is highly probable to be sampled around the value $\exp(\mu)$: then, an approximate relationship between $\xi^{(s+1)}$ and μ is derived, $\xi^{(s+1)} \approx \exp(\mu)$, or equivalently, $\mu \approx \log \xi^{(s+1)}$, which will be utilized shortly later. Throughout this paper, we use $\rho^2 = 0.001$ as the default value for ρ^2 .

Now, we describe how to calibrate μ by using the Hill estimator (15). We start with ordered realizations of $\lambda^{(s+1)} = (\lambda_1^{(s+1)}, \dots, \lambda_p^{(s+1)})^\top$ by $\lambda_{(1)}^{(s+1)} \geq \dots \geq \lambda_{(p)}^{(s+1)}$. The Hill estimator based on $\lambda^{(s+1)}$ is then

$$\widehat{\xi}_k(\lambda^{(s+1)}) = \frac{1}{k-1} \sum_{j=1}^{k-1} \log \left(\frac{\lambda_{(j)}^{(s+1)}}{\lambda_{(k)}^{(s+1)}} \right), \quad \text{for } k_L \leq k \leq k_U, \quad (17)$$

where $k_L = \lfloor p/10 \rfloor$ and $k_U = \lfloor 9p/10 \rfloor$, with $\lfloor \cdot \rfloor$ is the floor function. In high-dimensional statistical modeling, the cardinality of the set $\{k_L, \dots, k_U\} = \{\lfloor p/10 \rfloor, \dots, \lfloor 9p/10 \rfloor\} \subset \{2, \dots, p\}$ is still large, approximately, $\lfloor 4p/5 \rfloor$, enough to maintain the consistency of the Hill estimator. Note that estimates in (17) depend on k , which needs to be automated. To eliminate dependency on k , first, we average out the Hill estimators (17) over k , and then use the relation $\mu \approx \log \xi^{(s+1)}$, to get:

$$\widehat{\mu}(\lambda^{(s+1)}) = \log \{ \widehat{\xi}(\lambda^{(s+1)}) \} = \log \left\{ \frac{1}{k_U - k_L + 1} \sum_{k=k_L}^{k_U} \widehat{\xi}_k(\lambda^{(s+1)}) \right\}. \quad (18)$$

Note that $\widehat{\mu}(\lambda^{(s+1)})$ changes at each iteration of the Gibbs sampler, and tuned by $\lambda^{(s+1)}$ through the Hill estimator. In other words, $\widehat{\mu}(\lambda^{(s+1)})$ can be regarded as a calibrated hyper-parameter adapted via the p -realizations $\lambda^{(s+1)}$. By replacing μ with $\widehat{\mu}(\lambda^{(s+1)})$ and substituting $\rho^2 = 0.001$ in the full conditional posterior density $\pi(\xi| -)$ (16), the Gibbs sampler is now automated.

Finally, we explain how to sample from the density $\pi(\xi|-)$ (16). For that, first, use a change of variable $\eta = \log \xi$ and sample from

$$\eta^{(s+1)} \sim \pi(\eta|-) = \pi(\eta|\lambda^{(s+1)}, \tau^{(s+1)}) \propto \mathcal{L}(\eta) \cdot \mathcal{N}_1(\eta|\widehat{\mu}(\lambda^{(s+1)}), \rho^2 = 0.001), \quad (19)$$

where $\mathcal{L}(\eta) = \mathcal{V}_p(e^\eta)\mathcal{I}_{(\log 1/2, \infty)}(\eta) = [\{\Gamma(p/e^\eta + 1)\}^{-1}\pi^{p/2} \prod_{j=1}^p (\tau^{(s+1)} + e^\eta \lambda_j^{(s+1)})^{-(1/e^\eta+1)}] \mathcal{I}_{(\log 1/2, \infty)}(\eta)$. Once we obtain a sample $\eta^{(s+1)} \sim \pi(\eta|-)$, then $\xi^{(s+1)} \sim \pi(\xi|-)$ is obtained via the inverse transformation through $\xi^{(s+1)} = \exp \eta^{(s+1)}$.

We use the elliptical slice sampler (ESS) [21] to sample from $\eta^{(s+1)} \sim \pi(\eta|-)$ (19) that exploits the Gaussian prior measure. Conceptually, ESS and the Metropolis-Hastings (MH) algorithm are similar: both methods are comprised of two steps: *proposal step* and *criterion step*. A difference between the two algorithms arises in the *criterion step*. If the new candidate does not pass the criterion, then MH takes the current state as the next state: whereas, ESS re-proposes a new candidate until rejection does not take place, rendering the algorithm rejection-free. Further information for ESS is referred to the original paper [21]. Using a jargon from their paper, the calibrated μ , $\widehat{\mu}(\lambda^{(s+1)})$, is positioned at the center of an ellipse [21, 23]. Hence we refer to Algorithm 1 for the *Step 5* as the *elliptical slice sampler centered by the Hill estimator*.

Algorithm 1: Elliptical slice sampler centered by the Hill estimator

Circumstance : At the *Step 5* of the s -th iteration of the Gibbs sampler in Subsection 5.1.

Input : Current state $\xi^{(s)}$, and posterior realizations $\lambda^{(s+1)}$ and $\tau^{(s+1)}$ obtained from the *Step 3* and *Step 4*, respectively.

Output : A new state $\xi^{(s+1)}$.

1. Calibration of μ : obtain $\widehat{\mu}(\lambda^{(s+1)}) = \log \{\widehat{\xi}(\lambda^{(s+1)})\}$ (18).

2. Variable change ($\eta = \log \xi$): $\eta^{(s)} = \log \xi^{(s)}$.

3. Implement elliptical slice sampler to (19);

a. Choose ellipse centered by the Hill estimator: $\nu \sim \mathcal{N}_1(\widehat{\mu}(\lambda^{(s+1)}), \rho^2 = 0.001)$.

b. Define a criterion function:

$$\alpha(\eta, \eta^{(s)}) = \min\{\mathcal{L}(\eta)/\mathcal{L}(\eta^{(s)}), 1\} : (\log 1/2, \infty) \rightarrow [0, 1],$$

$$\text{where } \mathcal{L}(\eta) = [\{\Gamma(p/e^\eta + 1)\}^{-1} \pi^{p/2} \prod_{j=1}^p (\tau^{(s+1)} + e^\eta \lambda_j^{(s+1)})^{-(1/e^\eta + 1)}] \cdot \mathcal{I}_{(\log 1/2, \infty)}(\eta).$$

c. Choose a threshold and fix: $u \sim \mathcal{U}[0, 1]$.

d. Draw an initial proposal η^* :

$$\theta \sim \mathcal{U}(-\pi, \pi]$$

$$\eta^* = \{\eta^{(s)} - \widehat{\mu}(\lambda^{(s+1)})\} \cos \theta + \{\nu - \widehat{\mu}(\lambda^{(s+1)})\} \sin \theta + \widehat{\mu}(\lambda^{(s+1)})$$

e. **if** ($u < \alpha(\eta^*, \eta^{(s)})$) **{** $\eta^{(s+1)} = \eta^*$ **}** **else** **{**

Define a bracket : $(\theta_{\min}, \theta_{\max}] = (-\pi, \pi]$.

while ($u \geq \alpha(\eta^*, \eta^{(s)})$) **{**

Shrink the bracket and try a new point :

if ($\theta > 0$) $\theta_{\max} = \theta$ **else** $\theta_{\min} = \theta$

$$\theta \sim \mathcal{U}(\theta_{\min}, \theta_{\max}]$$

$$\eta^* = \{\eta^{(s)} - \widehat{\mu}(\lambda^{(s+1)})\} \cos \theta + \{\nu - \widehat{\mu}(\lambda^{(s+1)})\} \sin \theta + \widehat{\mu}(\lambda^{(s+1)})$$

}

$$\eta^{(s+1)} = \eta^*$$

}

4. Variable change ($\xi = e^\eta$): $\xi^{(s+1)} = \exp \eta^{(s+1)}$.

6. PROOF OF LEMMA 2

(a) Under the formulation (4) – (7), i.e., $\beta \sim \pi_{\text{GLT}}(\beta)$, we have

$$\begin{aligned}
\pi(\xi|-) &\propto \left\{ \prod_{j=1}^p \mathcal{GPD}(\lambda_j|\tau, \xi) \right\} \cdot \mathcal{IG}(\tau|p/\xi + 1, 1) \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi) \\
&= \left\{ \prod_{j=1}^p \frac{1}{\tau} \left(1 + \frac{\xi \lambda_j}{\tau} \right)^{-(1/\xi+1)} \right\} \cdot \frac{\tau^{-p/\xi-2} e^{-1/\tau}}{\Gamma(p/\xi + 1)} \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi) \\
&\propto \left\{ \tau^{p/\xi} \cdot \prod_{j=1}^p (\tau + \xi \lambda_j)^{-(1/\xi+1)} \right\} \cdot \frac{\tau^{-p/\xi-2}}{\Gamma(p/\xi + 1)} \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi) \\
&\propto \frac{\pi^{p/2}}{\Gamma(p/\xi + 1)} \prod_{j=1}^p (\tau + \xi \lambda_j)^{-(1/\xi+1)} \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) \cdot \mathcal{I}_{(1/2, \infty)}(\xi).
\end{aligned}$$

Now, our goal is to show

$$m(\lambda, \tau) = \int_{1/2}^{\infty} \frac{\pi^{p/2}}{\Gamma(p/\xi + 1)} \prod_{j=1}^p (\tau + \xi \lambda_j)^{-(1/\xi+1)} \cdot \log \mathcal{N}_1(\xi|\mu, \rho^2) d\xi < \infty, \quad \lambda \in (0, \infty)^p, \tau \in (0, \infty).$$

Let $x = 1/\xi$. Then

$$\begin{aligned}
m(\lambda, \tau) &= \int_2^0 \frac{\pi^{p/2}}{\Gamma(px + 1)} \prod_{j=1}^p \left(\frac{x}{\lambda_j + \tau x} \right)^{x+1} \cdot \log \mathcal{N}_1(1/x|\mu, \rho^2) \cdot -\frac{1}{x^2} dx \\
&= \pi^{p/2} \cdot \int_0^2 \frac{(1/\tau)^{p(x+1)}}{\Gamma(px + 1)} \prod_{j=1}^p \left(\frac{\tau x}{\lambda_j + \tau x} \right)^{x+1} \cdot \log \mathcal{N}_1(1/x|\mu, \rho^2) \cdot \frac{1}{x^2} dx \\
&\leq \pi^{p/2} \cdot \int_0^2 r(x) \cdot \log \mathcal{N}_1(1/x|\mu, \rho^2) \cdot \frac{1}{x^2} dx, \tag{20}
\end{aligned}$$

where $r(x) = (1/\tau)^{p(x+1)}/\Gamma(px + 1)$. Since $r(x)$ is continuous on a closed interval $[0, 2]$, there exists $x_0 \in [0, 2]$ such that $r(x_0) = \sup_{x \in [0, 2]} r(x) = B$. Using this bound B to (20), we have

$$\begin{aligned}
m(\lambda, \tau) &\leq \pi^{p/2} \cdot B \cdot \int_0^2 \log \mathcal{N}_1(1/x|\mu, \rho^2) \cdot \frac{1}{x^2} dx \\
&\leq \pi^{p/2} \cdot B \cdot \int_0^{\infty} \log \mathcal{N}_1(1/x|\mu, \rho^2) \cdot \frac{1}{x^2} dx = \pi^{p/2} \cdot B < \infty, \quad \lambda \in (0, \infty)^p, \tau \in (0, \infty).
\end{aligned}$$

(b) Start with a likelihood part:

$$\begin{aligned}
f(y|\xi) &= \int_0^\infty \int_0^\infty \int_{-\infty}^\infty \mathcal{N}_1(y|\beta, 1) \cdot \mathcal{N}_1(\beta|0, \lambda^2) \cdot \mathcal{GPD}(\lambda|\tau, \xi) \cdot \mathcal{IG}(1/\xi + 1, 1) d\beta d\lambda d\tau \\
&= \int_0^\infty \int_0^\infty \mathcal{N}_1(y|0, 1 + \lambda^2) \cdot \mathcal{GPD}(\lambda|\tau, \xi) \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\lambda d\tau \\
&= \frac{1}{\sqrt{2\pi}} \int_0^\infty \left(\int_0^\infty \frac{1}{\sqrt{1 + \lambda^2}} \cdot \exp \left\{ -\frac{y^2}{2(1 + \lambda^2)} \right\} \cdot \frac{1}{\tau} \left(1 + \frac{\xi\lambda}{\tau} \right)^{-(1/\xi+1)} d\lambda \right) \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\tau \\
&\leq \frac{1}{\sqrt{2\pi}} \int_0^\infty \left(\int_0^\infty \frac{1}{\sqrt{1 + \lambda^2}} \cdot \exp \left\{ -\frac{y^2}{2(1 + \lambda^2)} \right\} \cdot \frac{1}{\tau + \xi\lambda} d\lambda \right) \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\tau \\
&= \frac{1}{\sqrt{2\pi}} \int_0^\infty \left(\int_0^1 g(y, \lambda, \tau, \xi) d\lambda + \int_1^\infty g(y, \lambda, \tau, \xi) d\lambda \right) \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\tau, \quad (21)
\end{aligned}$$

where $g(y, \lambda, \tau, \xi) = \{1/\sqrt{1 + \lambda^2}\} \cdot \exp[-y^2/\{2(1 + \lambda^2)\}] \cdot \{1/(\tau + \xi\lambda)\}$, $y \in \mathfrak{R}$ and $\lambda, \tau > 0$. Because $g(y, \lambda, \tau, \xi)$ is continuous on a closed interval $[0, 1]$ as a function of λ , by mean value theorem for integral [2], there exists $c \in (0, 1)$ such that

$$\begin{aligned}
\int_0^1 g(y, \lambda, \tau, \xi) d\lambda &= g(y, c, \tau, \xi) = \frac{1}{\sqrt{1 + c^2}} \cdot \exp \left\{ -\frac{y^2}{2(1 + c^2)} \right\} \cdot \frac{1}{\tau + \xi c} \\
&\leq \left[\frac{1}{\sqrt{1 + c^2}} \exp \left\{ -\frac{y^2}{2(1 + c^2)} \right\} \right] \cdot \frac{1}{\tau} = A \cdot \frac{1}{\tau} \leq \frac{1}{\tau}, \quad \tau \in (0, \infty), \quad (22)
\end{aligned}$$

where $A = A(y, c) = \{1/(\sqrt{1 + c^2})\} \cdot \exp[-y^2/\{2(1 + c^2)\}]$, which is upper bounded by 1 on $\mathfrak{R} \times (0, 1)$. Also, we have

$$\begin{aligned}
\int_1^\infty g(y, \lambda, \tau, \xi) d\lambda &= \int_1^\infty \frac{1}{\sqrt{1 + \lambda^2}} \cdot \exp \left\{ -\frac{y^2}{2(1 + \lambda^2)} \right\} \cdot \frac{1}{\tau + \xi\lambda} d\lambda \\
&\leq \int_1^\infty \frac{1}{\lambda} \cdot 1 \cdot \frac{1}{\xi\lambda} d\lambda = \int_1^\infty \frac{1}{\lambda^2} d\lambda \cdot \frac{1}{\xi} = \frac{1}{\xi}, \quad \xi \in (1/2, \infty). \quad (23)
\end{aligned}$$

Using the upper bounds (22) and (23) to (21), then we have

$$\begin{aligned}
f(y|\xi) &\leq \frac{1}{\sqrt{2\pi}} \int_0^\infty \left(\frac{1}{\tau} + \frac{1}{\xi} \right) \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\tau = \frac{1}{\sqrt{2\pi}} \left(\int_0^\infty \frac{1}{\tau} \cdot \mathcal{IG}(\tau|1/\xi + 1, 1) d\tau + \frac{1}{\xi} \right) \\
&= \frac{1}{\sqrt{2\pi}} \left\{ \left(\frac{1}{\xi} + 1 \right) + \frac{1}{\xi} \right\} = \frac{1}{\sqrt{2\pi}} \left(\frac{2}{\xi} + 1 \right) \leq \frac{5}{\sqrt{2\pi}} < \infty, \quad y \in \mathfrak{R}, \xi \in (1/2, \infty).
\end{aligned}$$

Therefore, trivially for any proper prior $\pi(\xi)$ on $(1/2, \infty)$, we have

$$m(y) = \int_{1/2}^{\infty} f(y|\xi) \cdot \pi(\xi) d\xi \leq \frac{5}{\sqrt{2\pi}} \int_{1/2}^{\infty} \pi(\xi) d\xi = \frac{5}{\sqrt{2\pi}} < \infty, \quad y \in \mathfrak{R}.$$

7. PROOF OF PROPOSITION 2

Function $e^x E_1(x)$ satisfies tight upper and lower bounds [5];

$$\frac{1}{2} \cdot \log \left(\frac{x+2}{x} \right) < e^x E_1(x) < \log \left(\frac{x+1}{x} \right), \quad x > 0. \quad (24)$$

Replacing x with $Z_{\text{HS}}(\beta) = \beta^2/(2\tau^2)$ and multiplying $K_{\text{HS}} = 1/(\tau 2^{1/2} \pi^{3/2})$ to the both sides of the inequalities (24) lead to;

$$l(\beta) < \pi_{\text{HS}}(\beta|\tau) < u(\beta), \quad \beta \in \mathfrak{R}, \quad \tau > 0, \quad (25)$$

where $l(\beta) = (K_{\text{HS}}/2) \cdot \log \{(Z_{\text{HS}}(\beta) + 2)/Z_{\text{HS}}(\beta)\}$ and $u(\beta) = K_{\text{HS}} \cdot \log \{(Z_{\text{HS}}(\beta) + 1)/Z_{\text{HS}}(\beta)\}$.

Now, denote the tail (survival) function of the random variable $\beta|\tau$ given $\tau > 0$, by $\bar{F}_{\text{HS}}(\beta|\tau) = 1 - F_{\text{HS}}(\beta|\tau)$: that is, $(d/d\beta)F_{\text{HS}}(\beta|\tau) = \pi_{\text{HS}}(\beta|\tau)$. Then to show that the tail-index of $\pi_{\text{HS}}(\beta|\tau)$ is $\alpha = 1$ for any $\tau > 0$, it is sufficient to prove that $\lim_{\beta \rightarrow \infty} \bar{F}_{\text{HS}}(c\beta|\tau)/\bar{F}_{\text{HS}}(\beta|\tau) = c^{-1}$ for any $c > 0$ and $\tau > 0$ because $\pi_{\text{HS}}(\beta|\tau)$ is a symmetric density.

Using L'Hôpital's Rule, we have $\lim_{\beta \rightarrow \infty} \bar{F}_{\text{HS}}(c\beta|\tau)/\bar{F}_{\text{HS}}(\beta|\tau) = c \cdot \lim_{\beta \rightarrow \infty} \pi_{\text{HS}}(c\beta|\tau)/\pi_{\text{HS}}(\beta|\tau)$.

Now, use inequality (25) to bound the function inside of $\lim_{\beta \rightarrow \infty} \pi_{\text{HS}}(c\beta|\tau)/\pi_{\text{HS}}(\beta|\tau)$;

$$\frac{l(c\beta)}{u(\beta)} < \frac{\pi_{\text{HS}}(c\beta|\tau)}{\pi_{\text{HS}}(\beta|\tau)} < \frac{u(c\beta)}{l(\beta)}, \quad c > 0, \beta \in \mathfrak{R}, \quad \tau > 0. \quad (26)$$

First, calculate the limit of the upper bound in the inequality (26) at infinity by using L'Hôpital's

Rule again;

$$\begin{aligned}
\lim_{\beta \rightarrow \infty} \frac{u(c\beta)}{l(\beta)} &= 2 \lim_{\beta \rightarrow \infty} \frac{\log \{(Z_{\text{HS}}(c\beta) + 1)/Z_{\text{HS}}(c\beta)\}}{\log \{(Z_{\text{HS}}(\beta) + 2)/Z_{\text{HS}}(\beta)\}} \\
&= 2 \lim_{\beta \rightarrow \infty} \frac{\{Z_{\text{HS}}(c\beta)/(Z_{\text{HS}}(c\beta) + 1)\} \cdot (-c^2/Z_{\text{HS}}(c\beta)^2)}{\{Z_{\text{HS}}(\beta)/(Z_{\text{HS}}(\beta) + 2)\} \cdot (-2/Z_{\text{HS}}(\beta)^2)} \\
&= c^2 \cdot \lim_{\beta \rightarrow \infty} \frac{Z_{\text{HS}}(\beta) \cdot (Z_{\text{HS}}(\beta) + 2)}{Z_{\text{HS}}(c\beta) \cdot (Z_{\text{HS}}(c\beta) + 1)} = c^{-2}, \quad c > 0.
\end{aligned}$$

By the same way, we can show $\lim_{\beta \rightarrow \infty} l(c\beta)/u(\beta) = c^{-2}$, $c > 0$. Use the squeeze theorem to the inequality (26) to finish the proof.

8. PROOF OF PROPOSITION 3

(a) Clearly,

$$\pi(\beta|\tau, \xi) = \frac{1}{\tau\sqrt{2\pi}} \int_0^\infty \frac{1}{\lambda} \exp\left(-\frac{\beta^2}{2\lambda^2}\right) \left(1 + \frac{\xi\lambda}{\tau}\right)^{-(1/\xi+1)} d\lambda.$$

Let $x = \xi\lambda/\tau$. Then

$$\pi(\beta|\tau, \xi) = \frac{1}{\tau\sqrt{2\pi}} \int_0^\infty \exp\left(-\frac{\beta^2\xi^2}{2\tau^2x^2}\right) x^{-1}(1+x)^{-(1/\xi+1)} dx,$$

or equivalently, for $t = 1/x^2$:

$$\pi(\beta|\tau, \xi) = K \int_0^\infty e^{-Zt} (t^{1/2})^{-1+1/\xi} (1+t^{1/2})^{-(1+1/\xi)} dt, \tag{27}$$

where $K = 1/(\tau 2^{3/2} \pi^{1/2})$ and $Z(\beta) = \beta^2 \xi^2 / (2\tau^2)$. Use $Z = Z(\beta)$ to avoid notation clutter. To utilize the Newton's generalized binomial theorem;

$$(x+y)^r = \sum_{k=0}^{\infty} \binom{r}{k} x^{r-k} y^k, \quad |x| > |y|, r \in \mathcal{C},$$

we divide the integral in (27) into two parts. Then we have

$$\pi(\beta|\tau, \xi) = K \left\{ \int_0^1 e^{-Zt} (t^{1/2})^{-1+1/\xi} (1+t^{1/2})^{-(1+1/\xi)} dt + \int_1^\infty e^{-Zt} (t^{1/2})^{-1+1/\xi} (1+t^{1/2})^{-(1+1/\xi)} dt \right\}. \quad (28)$$

The first integral of (28) is

$$\begin{aligned} \int_0^1 e^{-Zt} (t^{1/2})^{-1+1/\xi} (1+t^{1/2})^{-(1+1/\xi)} dt &= \int_0^1 e^{-Zt} (t^{1/2})^{-1+1/\xi} \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} (t^{1/2})^k dt \\ &= \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} \int_0^1 e^{-Zt} t^{(1+1/\xi+k)/2-1} dt \\ &= \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} Z^{-(1+1/\xi+k)/2} \gamma\{(1+1/\xi+k)/2, Z\}, \end{aligned} \quad (29)$$

where $\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt$ ($s, x \in \mathfrak{R}$), is the incomplete lower gamma function.

The second integral of (28) is

$$\begin{aligned} \int_1^\infty e^{-Zt} (t^{1/2})^{-1+1/\xi} (1+t^{1/2})^{-(1+1/\xi)} dt &= \int_1^\infty e^{-Zt} (t^{1/2})^{-1+1/\xi} \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} (t^{1/2})^{-1-1/\xi-k} dt \\ &= \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} \int_0^1 e^{-Zt} t^{-1-k/2} dt \\ &= \sum_{k=0}^{\infty} \binom{-1-1/\xi}{k} E_{k/2+1}(Z), \end{aligned} \quad (30)$$

where $E_s(x) = \int_1^\infty e^{-xt} t^{-s} dt$ ($s, x \in \mathfrak{R}$) is the generalized exponential-integral function of real order [6, 20]. Use $\binom{-1-1/\xi}{k} = (-1)^k \binom{1/\xi+k}{k}$, (29), and (30) to conclude the proof.

(b) Prove by using the change of variable;

$$\begin{aligned} \pi(\kappa|\tau, \xi) = \mathcal{GPD}(\lambda|\tau, \xi) \Big|_{\lambda=\sqrt{(1-\kappa)/\kappa}} &\cdot \left| \frac{d\lambda}{dk} \right| = \frac{1}{\tau} \left(1 + \frac{\xi}{\tau} \sqrt{\frac{1-\kappa}{\kappa}} \right)^{-(1/\xi+1)} \cdot \frac{1}{2\kappa^2} \left(\frac{1-\kappa}{\kappa} \right)^{-1/2} \\ &= \frac{1}{2\tau} (\tau\sqrt{\kappa} + \xi\sqrt{1-\kappa})^{-(1/\xi+1)} (\tau\sqrt{\kappa})^{1/\xi+1} \cdot \frac{1}{\kappa^2} \left(\frac{1-\kappa}{\kappa} \right)^{-1/2} \\ &= \frac{\tau^{1/\xi}}{2} \cdot \frac{\kappa^{1/(2\xi)-1} (1-\kappa)^{-1/2}}{\{\tau\kappa^{1/2} + \xi(1-\kappa)^{1/2}\}^{(1+1/\xi)}}. \end{aligned}$$

9. PROOF OF COROLLARY 4

(a) In general, the generalized exponential-integral function has the following property; $\lim_{x \rightarrow 0^+} E_1(x) = \infty$ and $\lim_{x \rightarrow 0^+} E_s(x) = 1/(s-1)$ for $s > 1$ [6]. Using this property, if $k = 0$, then $\lim_{|\beta| \rightarrow 0} \psi_{k=0}^S(\beta) = \lim_{|\beta| \rightarrow 0} E_1\{Z(\beta)\} = \infty$ because $Z(\beta) = \beta^2 \xi^2 / (2\tau^2)$. If $k \in \{1, 2, \dots\}$, then $\lim_{\beta \rightarrow 0^+} \psi_k^S(\beta) = \lim_{\beta \rightarrow 0^+} E_{k/2+1}\{Z(\beta)\} = 2/k < \infty$.

(b) In general, the incomplete gamma function has the following property; $\lim_{x \rightarrow 0^+} x^{-a} \cdot \gamma(a, x) = a^{-1}$ for $a > 0$ [13]. Using this property, $\lim_{|\beta| \rightarrow 0} \psi_k^R(\beta) = \lim_{|\beta| \rightarrow 0} Z(\beta)^{-(1+1/\xi+k)/2} \cdot \gamma\{(1 + 1/\xi + k)/2, Z(\beta)\} = 2/(1 + 1/\xi + k) < \infty$ for all $k \in \{0\} \cup \{1, 2, \dots\}$.

(c) In general, the generalized exponential-integral function has the following property; $e^{-x}/(x + s) \leq E_s(x) \leq e^{-x}/(x + s - 1)$ for $x > 0$ and $s \geq 1$ [6]. Using this property, we obtain an inequality $e^{-Z(\beta)}/\{Z(\beta) + s\} \leq E_s(Z(\beta)) \leq e^{-Z(\beta)}/\{Z(\beta) + s - 1\}$ for $|\beta| > 0$ and $s \geq 1$. As $|\beta| \rightarrow \infty$, both bounds of $E_s(Z(\beta))$ converges to zero with squared exponential rate, and hence, $E_s(Z(\beta))$ also do for any $s \geq 1$.

(d) For fixed $k \in \{0\} \cup \{1, 2, \dots\}$ and ξ , we have $\lim_{|\beta| \rightarrow \infty} \gamma\{(1 + 1/\xi + k)/2, Z(\beta)\} = \Gamma((1 + 1/\xi + k)/2)$, where Γ is the gamma function, and hence, the function $\gamma\{(1 + 1/\xi + k)/2, Z(\beta)\}$ is a slowly varying function [19]. Using this we can re-express $\psi_k^R(\beta) = Z(\beta)^{-(1+1/\xi+k)/2} \cdot \gamma\{(1 + 1/\xi + k)/2, Z(\beta)\}$ by $\psi_k^R(\beta) = \beta^{-(1+1/\xi+k)} \cdot L(\beta)$, where L is a slowly varying function. This implies that the tail-index of function $\psi_k^R(\beta)$ is $1 + 1/\xi + k$.

REFERENCES

- [1] BAI, G., CHEUNG, I., SHULHA, H. P., COELHO, J. E., LI, P., DONG, X., JAKOVCEVSKI, M., WANG, Y., GRIGORENKO, A., JIANG, Y. et al. (2014). Epigenetic dysregulation of hairy and enhancer of split 4 (hes4) is associated with striatal degeneration in postmortem huntington brains. *Human molecular genetics* **24**, 1441–1456.
- [2] BARTLE, R. G. & SHERBERT, D. R. (2011). *Introduction to real analysis*. Hoboken, NJ: Wiley.

- [3] BHADRA, A., DATTA, J., POLSON, N. G., WILLARD, B. et al. (2017). The horseshoe+ estimator of ultra-sparse signals. *Bayesian Analysis* **12**, 1105–1131.
- [4] BISHOP, C. M. (2006). *Pattern recognition and machine learning*. springer.
- [5] CARVALHO, C. M., POLSON, N. G. & SCOTT, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika* **97**, 465–480.
- [6] CHICCOLI, C., LORENZUTTA, S. & MAINO, G. (1992). Concerning some integrals of the generalized exponential-integral function. *Computers & Mathematics with Applications* **23**, 13–21.
- [7] DAMLEN, P., WAKEFIELD, J. & WALKER, S. (1999). Gibbs sampling for bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **61**, 331–344.
- [8] DREES, H., DE HAAN, L., RESNICK, S. et al. (2000). How to make a hill plot. *The Annals of Statistics* **28**, 254–274.
- [9] DU, Y.-J., LUO, X.-Y., HAO, Y.-Z., ZHANG, T. & HOU, W.-R. (2007). cdna cloning and overexpression of acidic ribosomal phosphoprotein p1 gene (rplp1) from the giant panda. *International journal of biological sciences* **3**, 428.
- [10] EMBRECHTS, P., KLÜPPELBERG, C. & MIKOSCH, T. (2013). *Modelling extremal events: for insurance and finance*, vol. 33. Springer Science & Business Media.
- [11] GORRINGE, K. L., GEORGE, J., ANGLÉSIO, M. S., RAMAKRISHNA, M., ETEMAD-MOGHADAM, D., COWIN, P., SRIDHAR, A., WILLIAMS, L. H., BOYLE, S. E., YANAIHARA, N. et al. (2010). Copy number analysis identifies novel interactions between genomic loci in ovarian cancer. *PloS one* **5**, e11408.
- [12] HILL, B. M. (1975). A simple general approach to inference about the tail of a distribution. *The annals of statistics* , 1163–1174.

- [13] JAMESON, G. (2016). The incomplete gamma functions. *The Mathematical Gazette* **100**, 298–306.
- [14] JANSZ, N., KENIRY, A., TRUSSART, M., BILDSOE, H., BECK, T., TONKS, I. D., MOULD, A. W., HICKEY, P., BRESLIN, K., IMINITOFF, M. et al. (2018). Smchd1 regulates long-range chromatin interactions on the inactive x chromosome and at hox clusters. *Nature structural & molecular biology* **25**, 766.
- [15] KAR, S. P., BEESLEY, J., AL OLAMA, A. A., MICHAILIDOU, K., TYRER, J., KOTEJARAI, Z., LAWRENSEN, K., LINDSTROM, S., RAMUS, S. J., THOMPSON, D. J. et al. (2016). Genome-wide meta-analyses of breast, ovarian, and prostate cancer association studies identify multiple new susceptibility loci shared by at least two cancer types. *Cancer discovery* **6**, 1052–1067.
- [16] LEE, S. & KIM, J. H. (2018). Exponentiated generalized pareto distribution: Properties and applications towards extreme value theory. *Communications in Statistics-Theory and Methods* , 1–25.
- [17] LIU, J., LICHTENBERG, T., HOADLEY, K. A., POISSON, L. M., LAZAR, A. J., CHERNICK, A. D., KOVATICH, A. J., BENZ, C. C., LEVINE, D. A., LEE, A. V. et al. (2018). An integrated tcga pan-cancer clinical data resource to drive high-quality survival outcome analytics. *Cell* **173**, 400–416.
- [18] MAKALIC, E. & SCHMIDT, D. F. (2016). A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* **23**, 179–182.
- [19] MIKOSCH, T. (1999). *Regular variation, subexponentiality and their applications in probability theory*. Eindhoven University of Technology.
- [20] MILGRAM, M. (1985). The generalized integro-exponential function. *Mathematics of computation* **44**, 443–458.
- [21] MURRAY, I., PRESCOTT ADAMS, R. & MACKAY, D. J. (2010). Elliptical slice sampling .

- [22] NEAL, R. M. (2003). Slice sampling. *Annals of statistics* , 705–741.
- [23] NISHIHARA, R., MURRAY, I. & ADAMS, R. P. (2014). Parallel mcmc with generalized elliptical slice sampling. *The Journal of Machine Learning Research* **15**, 2087–2112.
- [24] PELLEGRINO, R., KAVAKLI, I. H., GOEL, N., CARDINALE, C. J., DINGES, D. F., KUNA, S. T., MAISLIN, G., VAN DONGEN, H. P., TUFIK, S., HOGENESCH, J. B. et al. (2014). A novel bhlhe41 variant is associated with short sleep and resistance to sleep deprivation in humans. *Sleep* **37**, 1327–1336.
- [25] PIIRONEN, J., VEHTARI, A. et al. (2017). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics* **11**, 5018–5051.
- [26] RESNICK, S. & STĂRICĂ, C. (1995). Consistency of hill’s estimator for dependent data. *Journal of Applied Probability* **32**, 139–167.
- [27] SHAW, J. A., PAGE, K., BLIGHE, K., HAVA, N., GUTTERY, D., WARD, B., BROWN, J., RUANGPRATHEEP, C., STEBBING, J., PAYNE, R. et al. (2012). Genomic analysis of circulating cell-free dna infers breast cancer dormancy. *Genome research* **22**, 220–231.
- [28] TIPPING, M. E. (2001). Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research* **1**, 211–244.
- [29] VAN DER PAS, S., KLEIJN, B., VAN DER VAART, A. et al. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics* **8**, 2585–2618.
- [30] VAN DER PAS, S. L., KLEIJN, B. J., VAN DER VAART, A. W. et al. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics* **8**, 2585–2618.
- [31] WANG, Q.-X., CHEN, E.-D., CAI, Y.-F., LI, Q., JIN, Y.-X., JIN, W.-X., WANG, Y.-H., ZHENG, Z.-C., XUE, L., WANG, O.-C. et al. (2016). A panel of four genes accurately

differentiates benign from malignant thyroid nodules. *Journal of Experimental & Clinical Cancer Research* **35**, 169.

[32] WASSERMAN, L. (2006). *All of nonparametric statistics*. Springer Science & Business Media.

[33] WU, D.-M., LIU, T., DENG, S.-H., HAN, R. & XU, Y. (2017). Slc39a4 expression is associated with enhanced cell migration, cisplatin resistance, and poor survival in non-small cell lung cancer. *Scientific reports* **7**, 7211.