

Shrinkage with Shrunk Shoulders: Gibbs Sampling Shrinkage Model Posteriors with Guaranteed Convergence Rates*

Akihiko Nishimura[†] and Marc A. Suchard[‡]

Abstract. Use of continuous shrinkage priors — with a “spike” near zero and heavy-tails towards infinity — is an increasingly popular approach to induce sparsity in parameter estimates. When the parameters are only weakly identified by the likelihood, however, the posterior may end up with tails as heavy as the prior, jeopardizing robustness of inference. A natural solution is to “shrink the shoulders” of a shrinkage prior by lightening up its tails beyond a reasonable parameter range, yielding a *regularized* version of the prior. We develop a regularization approach which, unlike previous proposals, preserves computationally attractive structures of original shrinkage priors. We study theoretical properties of the Gibbs sampler on resulting posterior distributions, with emphasis on convergence rates of the Pólya-Gamma Gibbs sampler for sparse logistic regression. Our analysis shows that the proposed regularization leads to geometric ergodicity under a broad range of global-local shrinkage priors. Essentially, the only requirement is for the prior $\pi_{\text{local}}(\cdot)$ on the local scale λ to satisfy $\pi_{\text{local}}(0) < \infty$. If $\pi_{\text{local}}(\cdot)$ further satisfies $\lim_{\lambda \rightarrow 0} \pi_{\text{local}}(\lambda)/\lambda^a < \infty$ for $a > 0$, as in the case of Bayesian bridge priors, we show the sampler to be uniformly ergodic.

Keywords: Bayesian inference, sparsity, generalized linear model, Markov chain Monte Carlo, ergodicity.

MSC2020 subject classifications: Primary 60J20, 62F15; secondary 62J07.

1 Introduction

Bayesian modelers are increasingly adopting continuous shrinkage priors to control the effective number of parameters and model complexity in a data-driven manner. These priors are designed to shrink most of the parameters towards zero while allowing for the likelihood to pull a small fraction of them away from zero. To achieve such effects, a shrinkage prior¹ has a density with a “spike” near zero and heavy-tails towards infinity, encoding information that parameter values are likely close to zero but otherwise could be anywhere. Originally developed for the purpose of sparse regression

*This work was partially supported through National Institutes of Health grants R01 AI107034, U19 AI135995 and R01 AI153044 and through Food and Drug Administration grant HHS 75F40120D00039.

[†]Department of Biostatistics, Johns Hopkins University Bloomberg School of Public Health, aki.nishimura@jhu.edu

[‡]Departments of Biomathematics, Biostatistics, and Human Genetics, University of California – Los Angeles, msuchard@ucla.edu

¹We drop the word “continuous” since “shrinkage priors” are commonly understood in the literature as continuous ones, which exclude traditional discrete spike-slab mixtures.

(Carvalho et al., 2009), shrinkage priors have found applications in trend filtering of time series data (Kowal et al., 2019), (dynamic) factor models (Kastner, 2019), graphical models (Li et al., 2019), compression of deep neural networks (Louizos et al., 2017), among others.

Of particular interest in this paper is an application of Bayesian shrinkage to a logistic regression model $y_i | \mathbf{x}_i, \boldsymbol{\beta} \sim \text{Bernoulli}(\text{logit}^{-1}(\mathbf{x}_i^\top \boldsymbol{\beta}))$ and computational properties of the corresponding posterior inference via Gibbs sampling. Due to the possibility of $\boldsymbol{\beta}$ being only weakly identifiable, use of a shrinkage prior on $\boldsymbol{\beta}$ here warrants proper modification of the prior’s tail in order to ensure reasonable computational and statistical behaviors. Under our tail regularization strategy, we show that the Gibbs sampler achieves geometric ergodicity under a broad range of shrinkage priors. Notably, our proof technique unifies analyses of the Gibbs samplers under various shrinkage priors, providing an easily verifiable condition for geometric and uniform ergodicity.

Shrinkage priors are often expressed as a scale mixture of Gaussians on the unknown parameter $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ (Polson and Scott, 2010):

$$\pi(\beta_j | \tau, \lambda_j) \sim \mathcal{N}(0, \tau^2 \lambda_j^2), \quad \lambda_j \sim \pi_{\text{loc}}(\cdot). \quad (1.1)$$

This *global-local* representation simplifies the posterior conditionals and leads to straightforward inference via Gibbs sampling. The *global scale* τ controls the average magnitude of β_j ’s and hence overall sparsity level. The *local scale* λ_j is specific to individual β_j and its density $\pi_{\text{loc}}(\cdot)$ controls the size of the spike and tail behavior of the marginal $\beta_j | \tau$. For instance, the popular *horseshoe* prior of Carvalho et al. (2010) uses $\pi_{\text{loc}}(\lambda) \propto (1 + \lambda^2)^{-1}$, inducing a marginal $\pi(\beta_j | \tau)$ with the spike proportional to $-\log(|\beta_j/\tau|)$ as $|\beta_j/\tau| \rightarrow 0$ and the tail proportional to $(\beta_j/\tau)^{-2}$ as $|\beta_j/\tau| \rightarrow \infty$. Another notable example is the Bayesian bridge prior of Polson et al. (2014), which generalizes the Bayesian lasso of Park and Casella (2008) with $\pi(\beta_j | \tau)$ having a larger spike as $|\beta_j/\tau| \rightarrow 0$ and heavier tails as $|\beta_j/\tau| \rightarrow \infty$. Most importantly from the computational efficiency perspective, the bridge prior possesses a closed-form expression $\pi(\beta_j | \tau) \propto \tau^{-1} \exp(-|\beta_j/\tau|^a)$ for $a \in (0, 1)$ and thus allows for a collapsed Gibbs update from $\tau | \boldsymbol{\beta}$ with λ_j ’s marginalized out.

For a simple purpose such as estimating the unknown means of independent Gaussian observations, a broad class of shrinkage priors achieve theoretically optimal performance (van der Pas et al., 2016; Ghosh and Chakrabarti, 2017). The lack of prior information in the tail of the distribution is problematic, however, in more complex models where parameters are only weakly identified. In such models, the posterior may have a tail as heavy as the prior, resulting in unreliable parameter estimates (Ghosh et al., 2018).

To address the above shortcoming of shrinkage priors, we build on the work of Piironen and Vehtari (2017) and propose a computationally convenient way to *regularize* shrinkage priors. The basic idea is to modify the prior so that the marginal distribution of $|\beta_j|$ has light-tails beyond a reasonable range. Our formulation has computational advantages over that of Piironen and Vehtari (2017) due to a subtle yet important difference. By preserving the global-local structure (1.1), our regularized shrinkage priors can benefit from partial marginalization approaches that substantially improve

mixing of Gibbs samplers (Polson et al. 2014; Johndrow et al. 2020; Nishimura and Suchard 2022, Appendix E). In addition, our regularization leaves the posterior conditionals of λ_j 's unchanged, allowing their conditional updates via existing specialized samplers (Griffin and Brown 2010; Polson et al. 2014; Nishimura and Suchard 2022, Appendix F).²

Our regularized shrinkage priors allow for posterior inference via Gibbs sampler whose convergence rates often are provably fast. As an illustrative example, we consider Bayesian sparse logistic regression models, whose need for regularization motivated the work of Piironen and Vehtari (2017). Gibbs sampling via the Pólya-Gamma data augmentation of Polson et al. (2013) is a state-of-the-art approach to posterior computation under logistic model. When combined with advanced numerical linear algebra techniques, this Gibbs sampler is highly scalable to large data sets (Nishimura and Suchard, 2022), but its theoretical convergence rate has not been investigated. Assuming that the prior density $\pi_{\text{loc}}(\lambda)$ is continuous and bounded except possibly at $\lambda = 0$, we establish that the Gibbs sampler is geometrically ergodic whenever $\pi_{\text{loc}}(0) < \infty$. Stronger uniform convergence is achieved when $\int \lambda^{-1} \pi_{\text{loc}}(\lambda) d\lambda < \infty$. The integrability condition holds in particular when $\pi_{\text{loc}}(\lambda) = O(\lambda^a)$ for $a > 0$ as $\lambda \rightarrow 0$, which is the case for normal-gamma priors with shape parameter larger than $1/2$ (Griffin and Brown, 2010) and for Bayesian bridge priors (Polson et al. 2014 and Nishimura and Suchard 2022, Appendix E).

Previous studies of the convergence rates under shrinkage models have focused exclusively on linear regression with specific parametric families of shrinkage priors (Pal and Khare, 2014; Johndrow et al., 2020). In contrast, our analysis requires no parametric assumptions on the shrinkage prior, at the same time extending the convergence results to the logistic model and, in Appendix A (Nishimura and Suchard, 2022), to the probit model.

To summarize, this work provides two major contributions to the Bayesian shrinkage literature. First, we propose an effective and Gibbs-friendly approach to suitably modify shrinkage priors for use in weakly-identifiable models (Section 2). Second, we develop theoretical tools to study the behavior of shrinkage model Gibbs samplers near the spike $\beta_j = 0$ without any parametric assumption on $\pi_{\text{loc}}(\cdot)$, thereby unifying convergence analyses of the logistic regression Gibbs samplers under a range of shrinkage priors (Section 3). We conclude the article in Section 4 by demonstrating a practical use case of regularized shrinkage models via simulation study, which emulates increasingly common situations where the sample sizes are large yet the signals are difficult to detect. Our simulation results in particular highlight the dual role of the regularization; by eliminating heavy-tails in the shrinkage model posterior, it induces both more stable parameter estimates and faster mixing of the Gibbs sampler.

²Appendix F (Nishimura and Suchard, 2022) describes a simple and provably efficient rejection-sampler for the conditional distributions of local scale parameter λ_j 's under the horseshoe prior. Despite the horseshoe's popularity, we find that no existing algorithm for the conditional update comes with theoretically guaranteed efficiency.

2 Regularized shrinkage prior

This section explains how our regularization approach allows us to incorporate prior information on the largest possible parameter values while maintaining the computational tractability of the original shrinkage prior.

Piironen and Vehtari (2017) proposes to control the tail behavior of a global-local shrinkage prior by defining its regularized version with *slab width* $\zeta > 0$ as

$$\beta_j | \tau, \lambda_j, \zeta \sim \mathcal{N}\left(0, \left(\frac{1}{\zeta^2} + \frac{1}{\tau^2 \lambda_j^2}\right)^{-1}\right), \quad (2.1)$$

with the prior $\pi_{\text{loc}}(\cdot)$ on the local scale λ_j unmodified. This regularization ensures that the variance of $\beta_j | \tau, \lambda_j, \zeta$ is upper bounded by ζ^2 and hence $\beta_j | \zeta$ marginally has a density with Gaussian tails beyond $|\beta_j| > \zeta$. The slab width ζ can be either given a prior distribution or fixed at a reasonable value.³

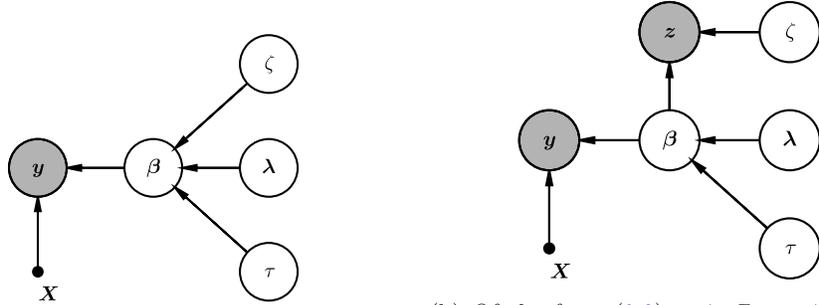
While beneficial in improving statistical properties (Piironen and Vehtari, 2017), regularization the form (2.1) compromises the posterior conditional structures of shrinkage models. Specifically, the conditional distribution of $\tau, \boldsymbol{\lambda}$ is altered through their dependency on ζ . This structural change is at best an inconvenience and potentially a cause of computational inefficiency, prohibiting the use of common acceleration techniques. For instance, the global scale τ is known to mix slowly when updating from its full conditional, so the state-of-the-art Gibbs samplers for Bayesian sparse regression marginalize out a subset of parameters when updating τ (Johndrow et al., 2020; Nishimura and Suchard, 2022). The analytical tractabilities of the integrals, which these marginalization strategies rely on, are lost when using the regularization as in (2.1).

We propose a more computationally convenient formulation, which induces regularization similar to that of (2.1) while keeping τ and $\boldsymbol{\lambda}$ conditionally independent of ζ given $\boldsymbol{\beta}$. Our regularized prior $\pi_{\text{reg}}(\cdot)$ defines the distribution of $\beta_j, \lambda_j | \tau, \zeta$ as

$$\begin{aligned} \pi_{\text{reg}}(\beta_j, \lambda_j | \tau, \zeta) &\propto \exp\left(-\frac{\beta_j^2}{2\zeta^2}\right) \frac{1}{\tau \lambda_j} \exp\left(-\frac{\beta_j^2}{2\tau^2 \lambda_j^2}\right) \pi_{\text{loc}}(\lambda_j) \\ &\propto \mathcal{N}\left(\beta_j \mid 0, \left(\frac{1}{\zeta^2} + \frac{1}{\tau^2 \lambda_j^2}\right)^{-1}\right) \left(1 + \frac{\tau^2 \lambda_j^2}{\zeta^2}\right)^{-1/2} \pi_{\text{loc}}(\lambda_j), \end{aligned} \quad (2.2)$$

where $\mathcal{N}(\cdot | 0, \sigma^2)$ denotes the centered Gaussian density with variance σ^2 . In other words, in addition to defining $\pi(\beta_j | \tau, \lambda_j, \zeta)$ as in (2.1), we alter the prior on λ_j as $\pi(\lambda_j | \tau, \zeta) \propto \pi_{\text{loc}}(\lambda_j) / \sqrt{1 + \tau^2 \lambda_j^2 / \zeta^2}$. Incidentally, we see that our regularized prior is very similar to that of Piironen and Vehtari (2017), but has a slightly lighter tail due to the factor $1 / \sqrt{1 + \tau^2 \lambda_j^2 / \zeta^2}$ which, as $\lambda_j \rightarrow \infty$, behaves like $\zeta / \tau \lambda_j$.

³While an appropriate choice of ζ is application specific, by way of illustration, we suggest $\zeta = 2$ as a weakly informative and sensible starting point in biomedical applications with standardized predictors. Schuemie et al. (2018) surveys 59,196 published effect estimates in the observational study literature and finds only a small portion of them exceeds 2.



(a) Of the form (2.1) as previously proposed. The posterior conditional of (τ, λ) is affected by their dependency on ζ through β .
 (b) Of the form (2.3) as in Proposition 2.1. Regularization does not affect the posterior conditional of (τ, λ) as the parameters remains decoupled from ζ .

Figure 2.1: Directed acyclic graphical model (a.k.a. Bayesian network) representation of regularized shrinkage priors under the two alternative formulations.

Alternatively, we can achieve the equivalent regularization through fictitious data that makes values $|\beta_j| \gg \zeta$ unlikely. While it may appear unnatural to introduce an auxiliary likelihood for the purpose of indirectly modifying a prior, this alternative formulation makes the regularization mechanism and resulting posterior properties more transparent. Figure 2.1 schematically describes this alternative construction of our regularized prior as well as the corresponding posterior structure when data \mathbf{y} and \mathbf{X} inform β through the likelihood $L(\mathbf{y} | \mathbf{X}, \beta)$.

Given a global-local prior $\beta_j | \tau, \lambda_j \sim \mathcal{N}(0, \tau^2 \lambda_j^2)$, we introduce fictitious data z_j whose realized value and underlying distribution are assumed to be

$$z_j = 0, \quad z_j | \beta_j, \zeta \sim \mathcal{N}(\beta_j, \zeta^2) \tag{2.3}$$

for $j = 1, \dots, p$. We then define the regularized prior as the distribution of β_j conditional on $z_j = 0$. Under this model, the distribution of $\beta_j | \tau, \lambda_j, \zeta, z_j = 0$ coincides with that of (2.1). On the other hand, the scale parameters τ, λ are conditionally independent of the others given β , so that the posterior full conditional $\tau, \lambda | \beta, \zeta, \mathbf{z}, \mathbf{y}, \mathbf{X}$ ($\stackrel{d}{=} \tau, \lambda | \beta$) has the same density as in the unregularized version. Our regularization thus allows the Gibbs sampler to update τ, λ with the exact same algorithm as the one designed for the original shrinkage prior. We summarize our discussion as Proposition 2.1 below.

Proposition 2.1. *Consider a global-local shrinkage prior $\beta_j | \tau, \lambda_j \sim \mathcal{N}(0, \tau^2 \lambda_j^2)$, $\lambda_j \sim \pi_{\text{loc}}(\cdot)$ and $\tau \sim \pi_{\text{glo}}(\cdot)$. Introducing the fictitious data $\mathbf{z} = \mathbf{0}$ as in (2.3) is equivalent to using the regularized prior (2.2) on (β_j, λ_j) , yielding*

$$\beta_j | \tau, \lambda_j, \zeta, z_j = 0 \sim \mathcal{N}\left(0, \left(\frac{1}{\zeta^2} + \frac{1}{\tau^2 \lambda_j^2}\right)^{-1}\right).$$

Or, with λ_j marginalized out, we have

$$\pi(\beta_j | \tau, \zeta, z_j = 0) \propto \pi(\beta_j | \tau) \exp\left(-\frac{\beta_j^2}{2\zeta^2}\right).$$

When the likelihood depends only on β , the posterior full conditional of τ, λ has density

$$\pi(\tau, \lambda | \beta) \propto \pi_{\text{glo}}(\tau) \prod_j \frac{1}{\tau \lambda_j} \exp\left(-\frac{\beta_j^2}{2\tau^2 \lambda_j^2}\right) \pi_{\text{loc}}(\lambda_j). \quad (2.4)$$

3 Geometric and uniform ergodicity under regularized sparse logistic regression

Shrinkage priors' popularity stems from, to a considerable extent, the ease of posterior computation via Gibbs sampling (Bhadra et al., 2019). As we have shown in Section 2, shrinkage models can incorporate regularization without affecting its computational tractability. We now investigate how fast such Gibbs samplers converge. While regularization was originally motivated to remedy statistically problematic behavior of heavy-tailed shrinkage priors, our results show that it can also improve the Gibbs samplers' convergence rates. The simulation results of Section 4 further corroborate the theory.

As a representative example where regularization is essential, we focus on Bayesian sparse logistic regression (Piironen and Vehtari, 2017; Nishimura and Suchard, 2022). To be explicit, we consider the model

$$\begin{aligned} y_i | \mathbf{x}_i, \beta &\sim \text{Bernoulli}(\text{logit}^{-1}(\mathbf{x}_i^\top \beta)), \\ z_j = 0 &\text{ for } z_j | \beta_j \sim \mathcal{N}(\beta_j, \zeta^2), \\ \beta_j | \tau, \lambda_j &\sim \mathcal{N}(0, \tau^2 \lambda_j^2), \tau \sim \pi_{\text{glo}}(\cdot), \lambda_j \sim \pi_{\text{loc}}(\cdot). \end{aligned} \quad (3.1)$$

The Pólya-Gamma data-augmentation of Polson et al. (2013) is a widely-used approach to carry out the posterior computation under the logistic model. By introducing an auxiliary parameter $\omega = (\omega_1, \dots, \omega_n)$ having a Pólya-Gamma distribution, the Gibbs sampler induces a transition kernel: $(\omega^*, \beta^*, \lambda^*, \tau^*) \rightarrow (\omega, \beta, \lambda, \tau)$ through the following cycle of conditional updates:

1. Draw $\tau | \beta^*, \lambda^*$ from the density proportional to (2.4). When using Bayesian bridge priors, draw from the collapsed distribution $\tau | \beta^*$ (Nishimura and Suchard 2022, Appendix E).
2. Draw $\lambda | \beta^*, \tau$ from the density proportional to (2.4).
3. Draw $\omega_i | \beta^*, \mathbf{X} \sim \text{PolyaGamma}(\text{shape} = 1, \text{tilting} = \mathbf{x}_i^\top \beta^*)$ for $i = 1, \dots, n$.
4. Draw $\beta | \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}$ from the multivariate-Gaussian

$$\begin{aligned} \beta | \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0} &\sim \mathcal{N}\left(\Phi^{-1} \mathbf{X}^\top (\mathbf{y} - \frac{1}{2}), \Phi^{-1}\right) \\ \text{for } \Phi &= \mathbf{X}^\top \Omega \mathbf{X} + \zeta^{-2} \mathbf{I} + \tau^{-2} \Lambda^{-2}, \end{aligned} \quad (3.2)$$

where $\Omega = \text{diag}(\omega)$ and $\Lambda = \text{diag}(\lambda)$.

Note that the transition kernel actually depends neither on ω^* nor τ^* (nor λ^* in the Bayesian bridge case) because of conditional independence. We refer readers to Polson et al. (2013) for more details on this data augmentation scheme. In our analysis, we do not use any specific properties of the Pólya-Gamma distribution aside from a couple of results from Choi and Hobert (2013) and Wang and Roy (2018).

The Pólya-Gamma Gibbs sampler for the logistic model has previously been analyzed under a Gaussian or flat prior on β (Choi and Hobert, 2013; Wang and Roy, 2018), but not under shrinkage priors. We establish geometric and uniform ergodicity — critical properties for any practical Markov chain Monte Carlo algorithms (Jones and Hobert, 2001). These properties imply the Markov chain central limit theorem and enables consistent estimation of Monte Carlo errors, ensuring that the Gibbs sampler reliably estimates quantities of interest (Flegal and Jones, 2011). To avoid cluttering notations and obscuring the main ideas, our analysis below assumes the slab width ζ to be fixed; however, the same conclusions hold if we only assume a prior constraint of the form $\zeta \leq \zeta_{\max} < \infty$ (Remark 3.9).

Below are the main ergodicity results we will establish in this section, the uniform rate under Bayesian bridge and geometric rate under more general shrinkage priors:

Theorem 3.1 (Uniform ergodicity in the Bayesian bridge case). *If the prior $\pi_{\text{glo}}(\cdot)$ is supported on $[\tau_{\min}, \infty)$ for $\tau_{\min} > 0$, then the Pólya-Gamma Gibbs sampler for regularized Bayesian bridge logistic regression is uniformly ergodic.*

Theorem 3.2 (Geometric ergodicity). *Suppose that the local scale prior satisfies $\|\pi_{\text{loc}}\|_{\infty} < \infty$ and that the global scale prior $\pi_{\text{glo}}(\cdot)$ is supported on $[\tau_{\min}, \tau_{\max}]$ for $0 < \tau_{\min} \leq \tau_{\max} < \infty$. Then the Pólya-Gamma Gibbs sampler for regularized sparse logistic regression is geometrically ergodic.*

Remark. Uniform/geometric ergodicity is an essential requirement for, yet not a guarantee of, practically efficient Markov chains (Roberts and Rosenthal, 2004). In fact, the simulation results of Section 4 show that the benefit of regularization is greatest when ζ is chosen small enough to impose a reasonable prior constraint on the value of β_j 's.

3.1 Proof approach: minorization and drift conditions

To establish Theorem 3.1 and 3.2, we verify that each Gibbs sampler satisfies the *minorization* and *drift* conditions, upon which geometric and uniform ergodicity are immediately implied by the well-known theory of Markov chains (Meyn and Tweedie, 2009; Roberts and Rosenthal, 2004). Here we introduce the relevant notions in terms of a generic transition kernel $P(\theta^*, d\theta)$.

In the statements to follow, we assume that $P(\theta^*, d\theta)$ has a corresponding density function which, with slight abuse of notation, we denote by $P(\theta | \theta^*)$; in other words, the two satisfy the relation $P(\theta^*, A) = \int_A P(\theta | \theta^*) d\theta$. A chain on the space $\theta \in \Theta$ with transition kernel $P(\theta^*, d\theta)$ is said to satisfy a minorization condition with a *small set* S if there are $\delta > 0$ and a probability density $\pi(\cdot)$ such that

$$P(\theta | \theta^*) \geq \delta \pi(\theta) \quad \text{for all } \theta^* \in S.$$

The chain is uniformly ergodic when $S = \Theta$. Otherwise, the chain is geometrically ergodic if it additionally satisfies a drift condition i.e. there is a *Lyapunov* function $V(\theta) \geq 0$ such that, for $\gamma < 1$ and $b < \infty$,

$$PV(\theta^*) := \int V(\theta)P(\theta | \theta^*) d\theta \leq \gamma V(\theta^*) + b$$

and $S = \{\theta : V(\theta) \leq d\}$ is a small set for some $d > 2b/(1 - \gamma)$ (Rosenthal, 1995).

For a two-block Gibbs sampler on the space (θ, ϕ) that alternately samples $\theta \sim \pi(\cdot | \phi)$ and $\phi \sim \pi(\cdot | \theta)$, geometric and uniform ergodicity of the joint chain follows from that of the marginal chain with transition kernel $P(\theta | \theta^*) = \int \pi(\theta | \phi)\pi(\phi | \theta^*) d\phi$ (Roberts and Rosenthal, 2001). In establishing the uniform ergodicity under the Bayesian bridge (Theorem 3.1), we decompose the collapsed Gibbs sampler into components β and (ω, τ, λ) and study the marginal chain in β . In the subsequent analysis establishing the geometric ergodicity under a more general class of regularized shrinkage priors (Theorem 3.2), we decompose the Gibbs sampler into components (β, λ) and (ω, τ) and study the marginal chain in (β, λ) .

3.2 Behavior of shrinkage model Gibbs samplers near $\beta_j = 0$

In many models, establishing minorization and drift condition amounts to quantifying the chain’s behavior in the tail of the target. In studying convergence rates under shrinkage models, however, we are faced with an additional and distinctive challenge: the need to establish that the chain does not get “stuck” near the spike at $\beta_j = 0$ (Pal and Khare, 2014; Johndrow et al., 2020). Regularization effectively eliminates the possibility of the chain meandering to infinity, making it relatively routine to analyze its behavior as $\beta_j \rightarrow \infty$. On the other hands, the existing results provide no general insights into the behavior near $\beta_j = 0$. In fact, a careful examination of the proofs by Pal and Khare (2014) and Johndrow et al. (2020) reveals that the analyses under various shrinkage priors could have been unified if we had a more general characterization of shrinkage model Gibbs samplers’ behavior near $\beta_j = 0$.

To fill in this theoretical gap, we start our analysis by abstracting key model-agnostic results from our proofs of minorization and drift condition for the sparse logistic regression Gibbs sampler. Our Propositions 3.3 and 3.4 below characterize properties of the distribution of $\lambda_j | \beta_j, \tau$ — this distribution, due to conditional independence, typically coincides with the full posterior conditional of λ_j and critically informs behavior of the subsequent update of β_j in a shrinkage model Gibbs sampler. Our proof techniques apply to a broad range of shrinkage priors, essentially requiring only that $\|\pi_{\text{loc}}\|_\infty := \max_\lambda \pi_{\text{loc}}(\lambda) < \infty$.⁴

Proposition 3.3 below plays a critical role in our proof of minorization condition. The proposition tells us that a sample from $\lambda_j | \beta_j^*, \tau$ has a uniformly lower-bounded probability of $\lambda_j \geq a$ as long as $|\beta_j^*/\tau|$ is bounded away from zero. In turn, the subsequent up-

⁴The results presented in this article, specifically those that depend on Proposition B.2 and Lemma B.3, implicitly assume that $\pi_{\text{loc}}(\lambda)$ is absolutely continuous at $\lambda_{\min} = \inf\{\lambda : \pi_{\text{loc}}(\lambda) > 0\}$. This is a purely technical assumption as any shrinkage prior in practice should satisfy $\pi_{\text{loc}}(\lambda) > 0$ for $\lambda > 0$ and be a differentiable function of λ .

date of β_j conditional on λ_j should also have a guaranteed chance of landing away from zero. Intuitively, we can thus interpret the proposition as suggesting that a shrinkage model Gibbs sampler should not get “absorbed” to the spike at $\beta_j = 0$. The difference in the limiting behavior as $|\beta_j^*/\tau| \rightarrow 0$, depending on whether $\int \lambda^{-1}\pi_{\text{loc}}(\lambda) d\lambda < \infty$, is also significant and leads to the difference between geometric and uniform convergence under the sparse logistic regression example through Theorem 3.6.

Proposition 3.3. *For any $a > 0$, the tail probability $\mathbb{P}(\lambda_j \geq a | \beta_j^*, \tau)$ is a decreasing function of $|\beta_j^*/\tau|$. If $\int \lambda^{-1}\pi_{\text{loc}}(\lambda) d\lambda = \infty$, then as $|\beta_j^*/\tau| \rightarrow 0$ the tail probability converges to 0, i.e. the conditional $\lambda_j | \beta_j^*, \tau$ converges in distribution to a delta measure at 0. If $\int \lambda^{-1}\pi_{\text{loc}}(\lambda) d\lambda < \infty$, then the conditional $\lambda_j | \beta_j^*, \tau$ converges in distribution to $\pi(\lambda_j) \propto \lambda_j^{-1}\pi_{\text{loc}}(\lambda_j)$ as $|\beta_j^*/\tau| \rightarrow 0$.*

Another key property of $\lambda_j | \beta_j, \tau$, featured prominently in our proof of the drift condition (Theorem 3.8), is provided by Proposition 3.4 below. To briefly provide a context, a Lyapunov function of the form $V(\beta) = \sum_j |\beta_j|^{-\alpha}$ has proven effective in analyzing a shrinkage model Gibbs sampler (Pal and Khare 2014, Johndrow et al. 2020, Section 3.4). And bounding the conditional expectation of $\tau^{-\alpha}\lambda_j^{-\alpha}$ as below often constitutes a critical step in establishing the drift condition.

Proposition 3.4. *Let $R > 0$ and $\alpha \in [0, 1)$. If $\|\pi_{\text{loc}}\|_\infty < \infty$, then there is an increasing function $\gamma(r) > 0$ with $\lim_{r \rightarrow 0} \gamma(r) = 0$, for which the expectation with respect to $\lambda_j | \beta_j^*, \tau$ satisfies*

$$\mathbb{E}[\tau^{-\alpha}\lambda_j^{-\alpha} | \tau, \beta_j^*] \leq \gamma(R/\tau) \left(|\beta_j^*|^{-\alpha} + |R|^{-\alpha} \right). \tag{3.3}$$

Proposition 3.3 and 3.4 are substantial theoretical contributions on their own, but we defer their proofs to Appendix B (Nishimura and Suchard, 2022) so that we can without interruption proceed to establish ergodicity results in the regularized sparse logistic regression case.

Remark. The assumption $\|\pi_{\text{loc}}\|_\infty < \infty$ is sufficient but not necessary one for the conclusion of Proposition 3.4 and later of Theorem 3.8. Following the analysis by Pal and Khare (2014), we can show that the conclusions also hold under normal-gamma priors with any shape parameter $a > 0$. These priors have the property $\pi_{\text{loc}}(\lambda) \sim O(\lambda^{2a-1})$ as $\lambda \rightarrow 0$ and hence $\lim_{\lambda \rightarrow 0} \pi(\lambda) = \infty$ for $a < 1/2$. We leave it as future work to characterize the behavior of general shrinkage priors with $\|\pi_{\text{loc}}\|_\infty = \infty$.

Remark. In Appendix A (Nishimura and Suchard, 2022), we show that Proposition 3.3 and 3.4 can also be applied to establish uniform/geometric ergodicity of a Gibbs sampler for Bayesian sparse probit regression, demonstrating their relevance beyond the sparse logistic regression example.

3.3 Minorization — with uniform ergodicity in special cases

Having described the noteworthy model-agnostic results within our proofs, from now on we focus exclusively on the regularized sparse logistic regression case. We first consider

the Gibbs sampler with fixed τ in Lemma 3.5 and Theorem 3.6. While fixing the global scale parameter is a common assumption in the ergodicity proofs for shrinkage models (Pal and Khare, 2014), we subsequently show that this assumption can be replaced with much weaker ones; we only require $\tau \sim \pi_{\text{glo}}(\cdot)$ to be supported away from 0 in Theorem 3.1 and additionally away from $+\infty$ in Theorem 3.7.

Let $P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda})$ denote the transition kernel corresponding to Step 3 and 4 of the Gibbs sampler as described in Page 372 and $P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau)$ corresponding to Step 2 – 4. In other words, we define

$$\begin{aligned} P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) &= \int \pi(\boldsymbol{\beta} | \boldsymbol{\omega}, \tau, \boldsymbol{\lambda}, \mathbf{y}, \mathbf{X}, z = \mathbf{0}) \pi(\boldsymbol{\omega} | \boldsymbol{\beta}^*, \mathbf{X}) d\boldsymbol{\omega}, \\ P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau) &= \int P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) \pi(\boldsymbol{\lambda} | \boldsymbol{\beta}^*) d\boldsymbol{\lambda}. \end{aligned}$$

The following lemma builds on a result of Choi and Hobert (2013) and plays a prominent role, along with Proposition 3.3, in our proofs of minorization conditions.

Lemma 3.5. *Whenever $\min_j \tau \lambda_j \geq R > 0$, there is $\delta' > 0$ — independent of τ and $\boldsymbol{\lambda}$ except through R — such that the following minorization condition holds:*

$$P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) \geq \delta' \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}),$$

where $\boldsymbol{\Phi}_R = \frac{1}{2} \mathbf{X}^\top \mathbf{X} + \zeta^{-2} \mathbf{I} + R^{-2} \mathbf{I}$ and $\boldsymbol{\mu}_R = \boldsymbol{\Phi}_R^{-1} \mathbf{X}^\top (\mathbf{y} - \mathbf{1}/2)$.

We defer the proof to Appendix C (Nishimura and Suchard, 2022).

We now establish a minorization condition for the Gibbs sampler with fixed τ .

Theorem 3.6 (Minorization). *Let $\epsilon, R > 0$. On a small set $\{\boldsymbol{\beta}^* : \min_j |\beta_j^*|/\tau \geq \epsilon\}$, the marginal transition kernel satisfies a minorization condition*

$$P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau) \geq \delta(\tau) \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}),$$

where $\boldsymbol{\mu}_R$ and $\boldsymbol{\Phi}_R$ are defined as in Lemma 3.5, and $\delta(\tau) > 0$ is increasing in τ and otherwise depends only on ϵ, R , and π_{loc} . Moreover, the minorization holds uniformly on $\boldsymbol{\beta}^* \in \mathbb{R}^p$ in case the prior satisfies $\int_0^\infty \lambda^{-1} \pi_{\text{loc}}(\lambda) d\lambda < \infty$.

Proof. Using Lemma 3.5, we have

$$\begin{aligned} P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau) &= \int P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) \pi(\boldsymbol{\lambda} | \boldsymbol{\beta}^*, \tau) d\boldsymbol{\lambda} \\ &\geq \int_{\{\min_j \tau \lambda_j \geq R\}} P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) \pi(\boldsymbol{\lambda} | \boldsymbol{\beta}^*, \tau) d\boldsymbol{\lambda} \\ &\geq \delta' \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}) \prod_j \int_{R/\tau}^\infty \pi(\lambda_j | \beta_j^*, \tau) d\lambda_j, \end{aligned}$$

for $\delta' > 0$ depending only on R . Also, Proposition 3.3 implies that whenever $|\beta_j^*|/\tau \geq \epsilon$

$$\int_{R/\tau}^\infty \pi(\lambda_j | \beta_j^*, \tau) d\lambda_j \geq \int_{R/\tau}^\infty \pi(\lambda | |\beta_j^*|/\tau = \epsilon) d\lambda > 0.$$

Hence, $\prod_j \int_{R/\tau}^\infty \pi(\lambda_j | \beta_j^*, \tau) d\lambda_j$ is lower bounded by a positive constant depending only on ϵ and R/τ . In case $C = \int_0^\infty \lambda^{-1} \pi_{\text{loc}}(\lambda) d\lambda < \infty$, we can forgo the assumption $|\beta_j^*/\tau| \geq \epsilon$ and obtain a uniform lower bound since

$$\int_R^\infty \pi(\lambda_j | \beta_j^*, \tau) d\lambda_j \geq \frac{1}{C} \int_R^\infty \lambda^{-1} \pi_{\text{loc}}(\lambda) d\lambda > 0. \quad \square$$

We now relax the assumption of fixed τ . The results of van der Pas et al. (2017) suggest that a constraint of the form $0 < \tau_{\min} \leq \tau \leq \tau_{\max} < \infty$ can improve the statistical property of shrinkage priors. As it turns out, such a constraint also enables us to establish minorization conditions for the full Gibbs sampler under sparse logistic regression with τ update incorporated. We can in fact take $\tau_{\max} = \infty$ in case of the Bayesian bridge prior, whose unique structure allows us to marginalize out λ_j 's when updating τ (Polson et al. 2014; Nishimura and Suchard 2022, Appendix E). This collapsed update of τ from $\tau | \beta$ makes it possible to deduce the uniform ergodicity result of Theorem 3.1 as an immediate consequence of Theorem 3.6 by studying the marginal transition $\beta^* \rightarrow \beta$ with kernel

$$P(\beta | \beta^*) = \int_{\tau_{\min}}^\infty P(\beta | \beta^*, \tau) \pi(\tau | \beta^*) d\tau. \quad (3.4)$$

Proof of Theorem 3.1. It suffices to establish uniform minorization for the marginal transition kernel (3.4). Under the Bayesian bridge prior, we have $\pi_{\text{loc}}(\lambda) \propto O(\lambda^{2a})$ as $\lambda \rightarrow 0$ (Nishimura and Suchard 2022, Appendix E) and hence $\int \lambda^{-1} \pi_{\text{loc}}(\lambda) < \infty$. The minorization condition of Theorem 3.6 thus holds uniformly in β^* , yielding

$$\int_{\tau_{\min}}^\infty P(\beta | \beta^*, \tau) \pi(\tau | \beta^*) d\tau \geq \mathcal{N}(\beta; \mu_R, \Phi_R^{-1}) \int_{\tau_{\min}}^\infty \delta(\tau) \pi(\tau | \beta^*) d\tau, \quad (3.5)$$

for $R > 0$. Theorem 3.6 further tells us that $\delta(\tau) > 0$ is increasing in τ , so we have

$$\int_{\tau_{\min}}^\infty \delta(\tau) \pi(\tau | \beta^*) d\tau \geq \delta(\tau_{\min}) > 0. \quad (3.6)$$

The inequalities (3.5) and (3.6) together establish uniform minorization. □

For more general shrinkage priors, the global scale τ must be updated from the full conditional $\tau | \beta, \lambda$. This makes it necessary to study the marginal transition $(\beta^*, \lambda^*) \rightarrow (\beta, \lambda)$, jointly in regression coefficients and local scales, with kernel

$$P(\beta, \lambda | \beta^*, \lambda^*) = \int_{\tau_{\min}}^{\tau_{\max}} P(\beta | \beta^*, \tau, \lambda) \prod_j \pi(\lambda_j | \beta_j^*, \tau) \pi(\tau | \beta^*, \lambda^*) d\tau. \quad (3.7)$$

We establish a minorization condition for this general case in Theorem 3.7.

Theorem 3.7. *If the prior $\pi_{\text{glo}}(\cdot)$ is supported on $[\tau_{\min}, \tau_{\max}]$ for $0 < \tau_{\min} \leq \tau_{\max} < \infty$, then the marginal transition kernel $P(\beta, \lambda | \beta^*, \lambda^*)$ of the Pólya-Gamma Gibbs sampler for regularized sparse logistic regression satisfies a minorization condition on a small set $\{(\beta^*, \lambda^*) : 0 < \epsilon \leq |\beta_j^*| \leq E < \infty \text{ for all } j\}$.*

Proof. By Lemma 3.5 and the fact $\tau\lambda_j \geq \tau_{\min}\lambda_j$, we know that for $R > 0$

$$P(\boldsymbol{\beta} | \boldsymbol{\beta}^*, \tau, \boldsymbol{\lambda}) \geq \mathbb{1}\{\min_j \tau_{\min}\lambda_j \geq R\} \delta' \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}). \quad (3.8)$$

To lower bound the term $\prod_j \pi(\lambda_j | \beta_j^*, \tau)$ in (3.7), we first recall that

$$\pi(\lambda_j | \beta_j^*, \tau) = \frac{\lambda_j^{-1} \exp(-\beta_j^{*2}/2\tau^2\lambda_j^2) \pi_{\text{loc}}(\lambda_j)}{\int_0^\infty \lambda^{-1} \exp(-\beta_j^{*2}/2\tau^2\lambda^2) \pi_{\text{loc}}(\lambda) d\lambda}.$$

When $\tau_{\min} \leq \tau \leq \tau_{\max}$ and $\epsilon \leq |\beta_j^*| \leq E$, we have

$$\exp(-E^2/2\tau_{\min}^2\lambda^2) \leq \exp(-\beta_j^{*2}/2\tau^2\lambda^2) \leq \exp(-\epsilon^2/2\tau_{\max}^2\lambda^2).$$

It follows from the above inequalities that

$$\pi(\lambda_j | \beta_j^*, \tau) \geq \frac{\lambda_j^{-1} \exp(-E^2/2\tau_{\min}^2\lambda_j^2) \pi_{\text{loc}}(\lambda_j)}{\int_0^\infty \lambda^{-1} \exp(-\epsilon^2/2\tau_{\max}^2\lambda^2) \pi_{\text{loc}}(\lambda) d\lambda} := \eta \pi_{\text{lower}}(\lambda_j) \quad (3.9)$$

for $\eta > 0$ and density $\pi_{\text{lower}}(\cdot)$ independent of β_j^* and τ . Combining (3.8) and (3.9), we can lower bound the transition kernel (3.7) as

$$\begin{aligned} P(\boldsymbol{\beta}, \boldsymbol{\lambda} | \boldsymbol{\beta}^*, \boldsymbol{\lambda}^*) &\geq \delta' \eta \mathbb{1}\left\{\min_j \lambda_j \geq \frac{R}{\tau_{\min}}\right\} \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}) \prod_j \pi_{\text{lower}}(\lambda_j) \int_{\tau_{\min}}^{\tau_{\max}} \pi(\tau | \boldsymbol{\beta}^*, \boldsymbol{\lambda}) d\tau \\ &= \delta' \eta \mathcal{N}(\boldsymbol{\beta}; \boldsymbol{\mu}_R, \boldsymbol{\Phi}_R^{-1}) \prod_j \mathbb{1}\left\{\lambda_j \geq \frac{R}{\tau_{\min}}\right\} \pi_{\text{lower}}(\lambda_j). \quad \square \end{aligned}$$

3.4 Drift condition and geometric ergodicity

Here we establish a drift condition for geometric ergodicity under sparse logistic regression. As discussed in Section 3.2, the regularization prevents the Markov chain from meandering to infinity, so the main question is whether the chain can get “stuck” for a long time near $\beta_j^* = 0$. The following result shows that this does not happen as long as the global scale τ is bounded away from zero.

Theorem 3.8. *Suppose that the local scale prior satisfies $\|\pi_{\text{loc}}\|_\infty < \infty$ and that the global scale prior $\pi_{\text{glo}}(\cdot)$ is supported on $[\tau_{\min}, \infty)$ for $\tau_{\min} > 0$. Then the marginal transition kernel $P(\boldsymbol{\beta}, \boldsymbol{\lambda} | \boldsymbol{\beta}^*, \boldsymbol{\lambda}^*)$ satisfies a drift condition with a Lyapunov function $V(\boldsymbol{\beta}) = \sum_j |\beta_j|^{-\alpha}$ for any $0 \leq \alpha < 1$.*

Proof. Note that $PV(\boldsymbol{\beta}^*)$ can be expressed as a series of iterated expectations with respect to (1) $\boldsymbol{\beta} | \boldsymbol{\omega}, \tau, \boldsymbol{\lambda}, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}$, (2) $\boldsymbol{\omega} | \boldsymbol{\beta}^*$, (3) $\boldsymbol{\lambda} | \boldsymbol{\beta}^*, \tau$, and (4) $\tau | \boldsymbol{\beta}^*, \boldsymbol{\lambda}^*$. We will bound the iterated expectations of $|\beta_j|^{-\alpha}$ one by one.

Since $\beta \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}$ is distributed as Gaussian, denoting by μ_j and σ_j^2 the conditional mean and variance of β_j , Proposition 3.10 below tells us that

$$\mathbb{E} [|\beta_j|^{-\alpha} \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}] \leq C_\alpha(\mu_j/\sigma_j) \sigma_j^{-\alpha},$$

$$\text{where } \sup_t C_\alpha(t) \leq \frac{\Gamma(\frac{1-\alpha}{2})}{2^{\alpha/2}\sqrt{\pi}} \text{ and } C_\alpha(t) = O(|t|^{-\alpha}) \text{ as } |t| \rightarrow \infty.$$

For the purpose of this proof, we can simply set C_α to be its global upper bound; however, a tighter bound may be obtained when the posterior concentrates away from zero and thereby resulting in $|\mu_j/\sigma_j| \rightarrow \infty$ and $C_\alpha(\mu_j/\sigma_j) \rightarrow 0$ as the sample size increases. Combined with Proposition 3.11 below, the above inequality implies

$$\frac{1}{C_\alpha} \mathbb{E} [|\beta_j|^{-\alpha} \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}] \leq \tau^{-\alpha} \lambda_j^{-\alpha} + \zeta^{-\alpha} + 1 - \frac{\alpha}{2} + \frac{\alpha}{2} \sum_{i=1}^n \omega_i x_{ij}^2. \tag{3.10}$$

In taking the expectation of (3.10) with respect to $\omega \mid \beta^*$, we use the result $\mathbb{E}[\omega_j \mid \beta^*] \leq 1/4$ of Wang and Roy (2018) to obtain

$$\frac{1}{C_\alpha} \mathbb{E} [|\beta_j|^{-\alpha} \mid \tau, \lambda] \leq \tau^{-\alpha} \lambda_j^{-\alpha} + \zeta^{-\alpha} + 1 - \frac{\alpha}{2} + \frac{\alpha}{8} \sum_{i=1}^n x_{ij}^2. \tag{3.11}$$

Taking the expectation of (3.11) with respect to $\lambda \mid \tau, \beta^*$, we have

$$\frac{1}{C_\alpha} \mathbb{E} [|\beta_j|^{-\alpha} \mid \tau, \beta^*] \leq \mathbb{E}[\tau^{-\alpha} \lambda_j^{-\alpha} \mid \tau, \beta_j^*] + C'(\alpha, \mathbf{X}),$$

$$\text{where } C'(\alpha, \mathbf{X}) = \zeta^{-\alpha} + 1 - \frac{\alpha}{2} + \frac{\alpha}{8} \sum_{i=1}^n x_{ij}^2. \tag{3.12}$$

Now choose $R > 0$ small enough that $\gamma(R/\tau) \leq \gamma(R/\tau_{\min}) < C_\alpha^{-1}$ in Proposition 3.4. Then we have the following inequality for $\gamma' := C_\alpha \gamma(R/\tau_{\min}) < 1$:

$$C_\alpha \mathbb{E}[\tau^{-\alpha} \lambda_j^{-\alpha} \mid \tau, \beta_j^*] \leq \gamma' (|\beta_j^*|^{-\alpha} + |R|^{-\alpha})$$

for all $\tau \geq \tau_{\min}$. Incorporating the above inequality into (3.12), we obtain

$$\mathbb{E} [|\beta_j|^{-\alpha} \mid \tau, \beta^*] \leq \gamma' |\beta_j^*|^{-\alpha} + \gamma' |R|^{-\alpha} + C_\alpha C'(\alpha, \mathbf{X}).$$

Since $\pi(\tau \mid \beta^*, \lambda^*)$ is supported on $\tau \geq \tau_{\min}$ by our assumption, taking the expectation with respect to $\tau \mid \beta^*, \lambda^*$ yield

$$\mathbb{E} [|\beta_j|^{-\alpha} \mid \beta^*, \lambda^*] \leq \gamma' |\beta_j^*|^{-\alpha} + \gamma' |R|^{-\alpha} + C_\alpha C'(\alpha, \mathbf{X}). \quad \square$$

Theorem 3.7 and 3.8 together imply the geometric ergodicity result of Theorem 3.2:

Proof of Theorem 3.2. We show that $V(\beta) = \sum_j |\beta_j|^{-\alpha} + \|\beta\|^2$ is a Lyapunov function for the marginal transition kernel $P(\beta, \lambda \mid \beta^*, \lambda^*)$. Note that

$$\begin{aligned} \mathbb{E} [\|\beta\|^2 \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}] &= \|\mathbb{E}[\beta \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}]\|^2 + \sum_j \text{var}(\beta_j^2 \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}) \\ &= \|\Sigma \mathbf{X}^\top (\mathbf{y} - \frac{1}{2})\|^2 + \sum_j \mathbf{e}_j^\top \Sigma \mathbf{e}_j \end{aligned}$$

for $\Sigma = (\mathbf{X}^\top \Omega \mathbf{X} + \zeta^{-2} \mathbf{I} + \tau^{-2} \Lambda^{-2})^{-1}$. Since $\Sigma \prec \zeta^2 \mathbf{I}$, we have $\mathbf{e}_j^\top \Sigma \mathbf{e}_j \leq \zeta^2$ and $\|\Sigma \mathbf{X}^\top (\mathbf{y} - \frac{1}{2})\|^2 \leq \zeta^2 \|\mathbf{X}^\top (\mathbf{y} - \frac{1}{2})\|^2$. Thus we have

$$\mathbb{E}[\|\beta\|^2 \mid \omega, \tau, \lambda, \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}] \leq \zeta^2 \|\Sigma \mathbf{X}^\top (\mathbf{y} - \frac{1}{2})\|^2 + n\zeta^2. \quad (3.13)$$

Since the right-hand side does not depend on ω, τ, λ , the expectation with respect to $P(\beta, \lambda \mid \beta^*, \lambda^*)$ satisfies the same bound:

$$\mathbb{E}[\|\beta\|^2 \mid \beta^*, \lambda^*] \leq \zeta^2 \|\Sigma \mathbf{X}^\top (\mathbf{y} - \frac{1}{2})\|^2 + n\zeta^2.$$

In addition to the above bound, we know that $\sum_j |\beta_j|^{-\alpha}$ is a Lyapunov function by Theorem 3.8. Hence, $V(\beta) = \sum_j |\beta_j|^{-\alpha} + \|\beta\|^2$ is again a Lyapunov function. Moreover, by Theorem 3.7, we know that the Gibbs sampler satisfies a minorization condition on the set $\{\beta^* : 0 < \epsilon \leq |\beta_j^*| \leq E < \infty \text{ for all } j\}$ for $\epsilon > 0$ and $E < \infty$. Thus the sampler is geometrically ergodic. \square

Remark 3.9. As mentioned earlier, the geometric and uniform ergodicity as well as analogues of the intermediate results continue to hold when we relax the assumption of fixed ζ to a prior constraint of the form $\zeta \leq \zeta_{\max} < \infty$. The proof goes as follows. Due to the conditional independence, the Gibbs sampler on the joint space draws alternately from $\zeta \mid \beta, \mathbf{z} = \mathbf{0}$ and $\beta, \omega, \tau, \lambda \mid \mathbf{y}, \mathbf{X}, \mathbf{z} = \mathbf{0}, \zeta$. By repeating all the previous arguments with ζ_{\max} in place of ζ , we obtain essentially the identical minorization and drift bounds that hold for all $\zeta \leq \zeta_{\max}$. Since the bounds hold uniformly on the support $\zeta \leq \zeta_{\max}$, the identical bounds again hold when taking the expectation over $\zeta \mid \beta, \mathbf{z} = \mathbf{0}$.

Auxiliary results for proof of geometric ergodicity

Proposition 3.10 and 3.11 below are used in the proof of Theorem 3.8 and are proved in Appendix D (Nishimura and Suchard, 2022). Proposition 3.10 is a refinement of Proposition A1 in Pal and Khare (2014) and of Equation (41) in Johndrow et al. (2020), neither of which have the $D(\mu/\sigma)$ term.

Proposition 3.10. *For $\alpha \in (0, 1)$ and $\beta \sim \mathcal{N}(\mu, \sigma^2)$, we have*

$$\mathbb{E}|\beta|^{-\alpha} \leq \frac{\Gamma(\frac{1-\alpha}{2})}{2^{\alpha/2} \sqrt{\pi}} \sigma^{-\alpha} \min\{1, D(\mu/\sigma)\},$$

where $D(t) = O(|t|^{-\alpha}) \rightarrow 0$ as $|t| \rightarrow \infty$ and can be chosen as

$$D(t) = \frac{1}{B(\frac{\alpha}{2}, \frac{1-\alpha}{2})} \left[\frac{2^{\frac{5}{2}-\alpha}}{1-\alpha} \exp\left(-\frac{t^2}{4}\right) + 2^{\frac{1}{2}+\alpha} \Gamma\left(\frac{\alpha}{2}\right) |t|^{-\alpha} \right]. \quad (3.14)$$

Proposition 3.11. *The diagonals σ_j of $\Sigma = (\mathbf{X}^\top \Omega \mathbf{X} + \zeta^{-2} \mathbf{I} + \tau^{-2} \Lambda^{-2})^{-1}$ satisfy the following inequality for $0 \leq \alpha < 1$:*

$$\sigma_j^{-\alpha} \leq \tau^{-\alpha} \lambda_j^{-\alpha} + \zeta^{-\alpha} + 1 - \frac{\alpha}{2} + \frac{\alpha}{2} \sum_{i=1}^n \omega_i x_{ij}^2.$$

4 Simulation

We run a simulation study to assess the computational and statistical properties of the regularized sparse logistic regression model. We use the Bayesian bridge prior $\pi(\beta_j | \tau) \propto \tau^{-1} \exp(-|\beta_j/\tau|^a)$ to take advantage of the efficient global scale parameter update scheme. This prior also allows us to experiment with a range of spike and tail behavior by varying the exponent a , inducing larger spikes and heavier tails as $a \rightarrow 0$. For the global scale parameter, we chose the objective prior $\pi_{\text{glo}}(\tau) \propto \tau^{-1}$ (Berger et al., 2015, Nishimura and Suchard 2022, Appendix E) with the range restriction $10^{-6} \leq \mathbb{E}[|\beta_j| | \tau] \leq 1$ to ensure posterior propriety, though in practice we never observe a posterior draw of τ outside this range. For the posterior computations, we use the Pólya-Gamma Gibbs sampler provided by the *bayesbridge* package available from Python Package Index (pypi.org); the source code is available at the GitHub repository <https://github.com/OHDSI/bayes-bridge>.

4.1 Data generating process: “large n , but weak signal” problems

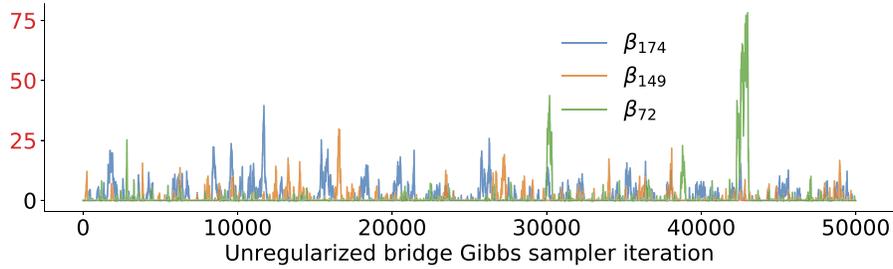
Piironen and Vehtari (2017) demonstrate the benefits of regularizing shrinkage priors in the “ $p > n$ ” case, when the number of predictors p exceeds the sample size n . To complement their study, we consider the case of rare outcomes and infrequently observed features, another common situation in which regularizing shrinkage priors becomes essential. For example in healthcare data, many outcomes of interests have low incidence rates and many treatments are prescribed to only a small fraction of patients (Tian et al., 2018). This results in binary outcomes \mathbf{y} and features \mathbf{x}_j filled mostly with 0’s, making the amount of information much less than otherwise expected (Greenland et al., 2016).

To simulate under these “large n , but weak signal” settings, we generate synthetic data with $n = 2,500$ and $p = 500$ as follows. We construct binary features with a range of observed frequencies by first drawing $2w_j \sim \text{Beta}(1/2, 2)$ for $j = 1, \dots, 500$; this in particular means $0 \leq w_j \leq 0.5$ and $\mathbb{E}[w_j] = 0.1$. For each j , we then generate $x_{ij} \sim \text{Bernoulli}(w_j)$ for $i = 1, \dots, n$. We choose the true signal to be $\beta_j = 1$ for $j = 1, \dots, 10$ and $\beta_j = 0$ for $j = 11, \dots, 500$. To simulate an outcome with low incidence rate, we choose the intercept to be $\beta_0 = 1.5$ and draw $y_i \sim \text{Bernoulli}(\text{logit}(-\mathbf{x}_i^\top \boldsymbol{\beta}))$, resulting in $y_i = 1$ for approximately 5% of its entries.

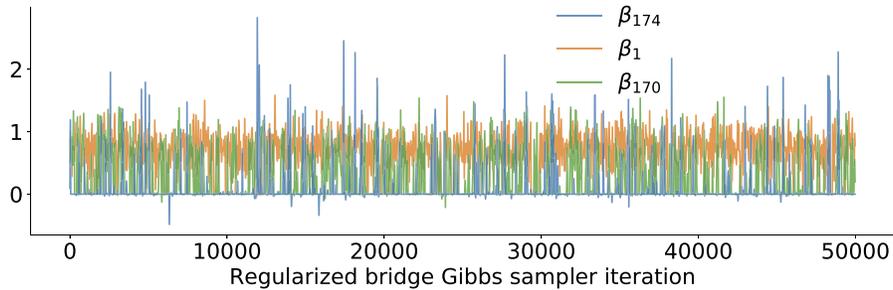
4.2 Convergence and mixing: with and without regularization

With the above data generating process, outcome \mathbf{y} and design matrix \mathbf{X} barely contain enough information to estimate all the coefficients β_j ’s. In particular, sparse logistic model without regularization can lead to a heavy-tailed posterior, for which uniform and geometric ergodicity of the Pólya-Gamma Gibbs sampler becomes questionable.

These potential convergence and mixing issues are evidenced by the traceplot (Figure 4.1a) of the posterior samples based on bridge exponent $a = 1/16$. As we are particularly concerned with the Markov chain wandering off to the tail of the target, we examine the estimated credible intervals to identify the coefficients with potential



(a) Without regularization, the Markov chain takes multiple “excursions” — each lasting over hundreds of iterations — into the unreasonable value range of the coefficients. The deviation in β_{72} is particularly prominent around the 42,000th iteration. More severe deviations may occur if the chain is run longer.



(b) With regularization, the Markov chain does not display any serious mixing issues. The noticeable auto-correlation is due to the multi-modality of the posterior, an unavoidable feature of shrinkage models. Note that the coefficients with widest credible intervals do not coincide with the unregularized setting.

Figure 4.1: Traceplot under the Bayesian bridge logistic regression with exponent $1/16$. Shown are the three coefficients with most potentially problematic mixing behaviors; see the main text for the details on our criteria.

convergence and mixing issues. Plotted in Figure 4.1 are the coefficients with the widest 95% credible intervals; these coefficients also have some of the smallest estimated effective sample sizes, though the accuracy of such estimates is not guaranteed without geometric ergodicity. When regularizing the shrinkage prior with a slab width $\zeta = 1$, the posterior samples indicate no such convergence or mixing issues (Figure 4.1b) and yield more sensible posterior credible intervals (Figure 4.2).

We emphasize that there is no fundamental change in the Gibbs sampler itself when incorporating the regularization, the only change being the addition of the $\zeta^{-2}\mathbf{I}$ term in the conditional precision matrix (3.2) when updating β . It is the change in the posterior — more specifically the guaranteed light tails of the β marginal — that induces faster convergence and mixing.

We also assess sensitivity of convergence and mixing rates on the slab width ζ . The regularized prior recovers the unregularized one as $\zeta \rightarrow \infty$. This means that, as seen

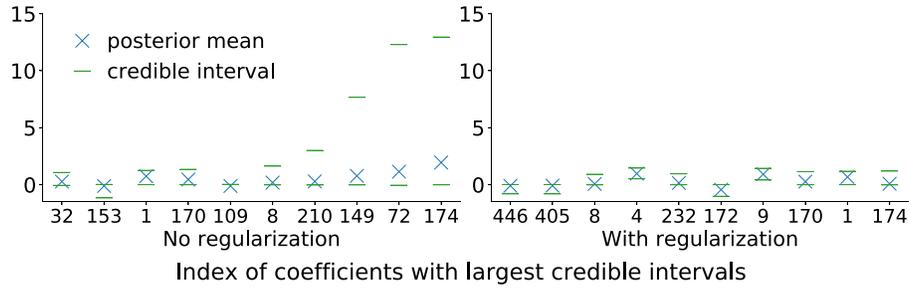


Figure 4.2: Ten widest 95% posterior credible intervals under the Bayesian bridge logistic regression with (right) and without (left) regularization. Without regularization, the intervals are unrealistically large compared to the signal size of $\beta_j = 1$ for $j = 1, \dots, 10$.

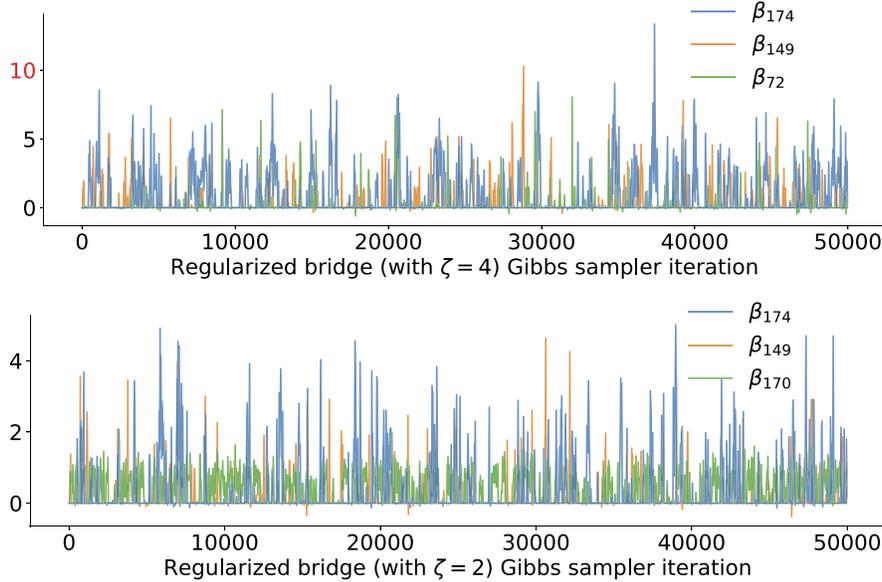


Figure 4.3: Traceplots under different slab widths: $\zeta = 2$ (bottom) and $\zeta = 4$ (top). The settings are otherwise identical to those of Figure 4.1. As before, the three coefficients with most problematic mixing behaviors do not always coincide across different slab widths.

from the problematic computational behavior of the unregularized model, ζ cannot be taken too large in this limited data setting. In other words, the choice of ζ has to reflect some degree of prior information on β_j 's. We need not assume strong prior information, however; Figure 4.3 demonstrates that even small amount of regularization (e.g. $\zeta = 2$ or 4) can noticeably improve the computational behavior over the unregularized case.

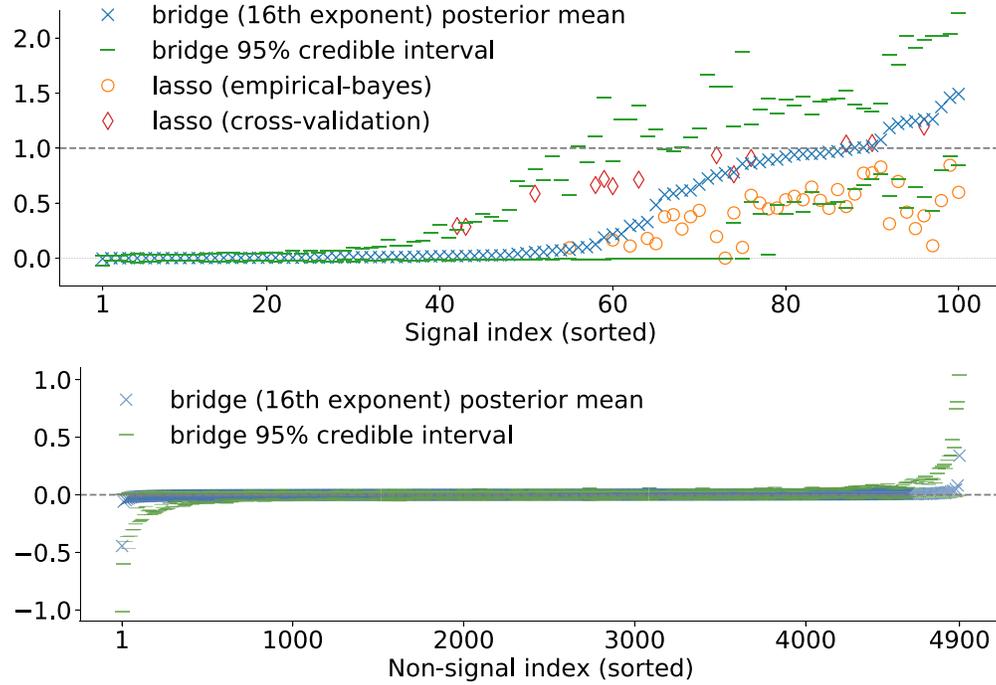


Figure 4.4: The 95% posterior credible intervals for the signals $\beta_j = 1$ (top) and non-signals $\beta_j = 0$ (bottom) under the Bayesian bridge logistic regression with the bridge exponent $1/16$. The intervals are sorted by the posterior means. To avoid clutter, the top plot shows only the non-zero values of the lasso estimates. The lasso estimates for the non-signals are summarized in Figure 4.5 and are not shown in the bottom plot.

4.3 Statistical properties of shrinkage model for weak signals

To study the shrinkage model's ability to separate out the non-zero β_j from the $\beta_j = 0$, we simulate 10 replicate data sets and estimate the posterior for each of them. In total, we obtain 5,000 marginal posterior distributions — 10 independent replications for each of the $p = 500$ regression coefficients — with 100 for the signal $\beta_j = 1$ and 4,900 for the non-signal $\beta_j = 0$. As all the predictors \mathbf{x}_j 's are simulated in an exchangeable manner, the 100 (and 4,900) posterior marginals for the signal (and non-signal) are also exchangeable.

Figure 4.4 show the posterior credible intervals. Due to the low incidence rate and infrequent binary features, many of the signals are too weak to be detected. We also find that the credible intervals seemingly do not achieve their nominal frequentist coverage for signals below detection strength. This finding is consistent with the existing theoretical results on shrinkage priors and is unsurprising in light of the impossibility theorem by Li (1989) — confidence intervals cannot be optimally tight and have nominal coverage at the same time. Credible intervals produced by Bayesian shrinkage models tend

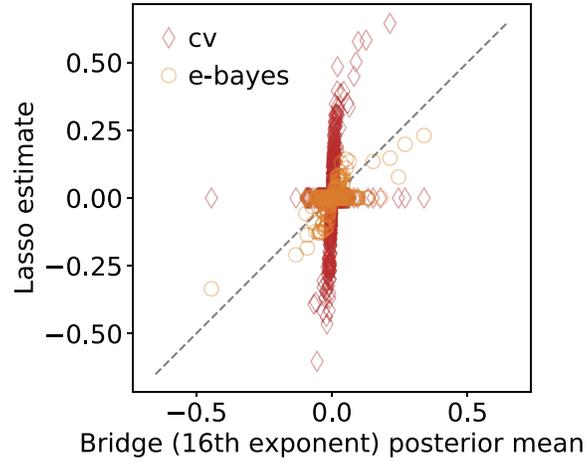


Figure 4.5: Comparison of the 4,900 Bayesian bridge posterior means and lasso estimates for the non-signals $\beta_j = 0$. Lasso with cross-validation produces a larger number of false positives. Lasso with the empirical Bayes calibration yields the estimates more in line with the bridge posterior.

to be optimally tight and thus require appropriate manual adjustments to achieve the nominal coverage (van der Pas et al., 2017). No statistical procedure is immune to this tightness-coverage trade-off; therefore, the apparent under-coverage should be seen not as a flaw but more as a feature of Bayesian shrinkage models.

We benchmark the signal detection capability of the posterior against the frequentist lasso, arguably the most widely-used approach to feature selection. Obtaining the lasso point estimates requires a selection of the hyper-parameter commonly referred to as the *penalty* parameter. For its choice, we first follow the standard practice of minimizing the ten-fold cross-validation errors (Hastie et al., 2009). However, this approach yields inconsistent and poor overall performance, detecting only 13 out of the 100 signals (Figure 4.4). Cross-validation likely fails here because each fold does not capture the characteristics of the whole data when the signals are so weak. As a more stable alternative for calibrating the penalty parameter, we try an empirical Bayes procedure based on the Bayesian interpretation of the lasso (Park and Casella, 2008). We first estimate the posterior marginal mean of the penalty parameter from the Bayesian lasso Gibbs sampler. Conditionally on this value, we then find the posterior mode of β . This procedure seems to yield more consistent performance, detecting 39 out of the 100 signals albeit with the estimates more shrunk towards null than the Bayesian posterior means. The empirical Bayes procedure demonstrates more consistent behavior for the non-signals as well (Figure 4.5).

We also assess how the spike size and (pre-regularization) tail behavior of the prior influence statistical properties of the resulting posterior. For this purpose, we fit the regularized bridge model with the exponent $a^{-1} \in \{2, 4, 8, 16\}$ to the same data sets. Figure 4.6 summarizes the credible intervals under the $a = 1/4$ case. The credible

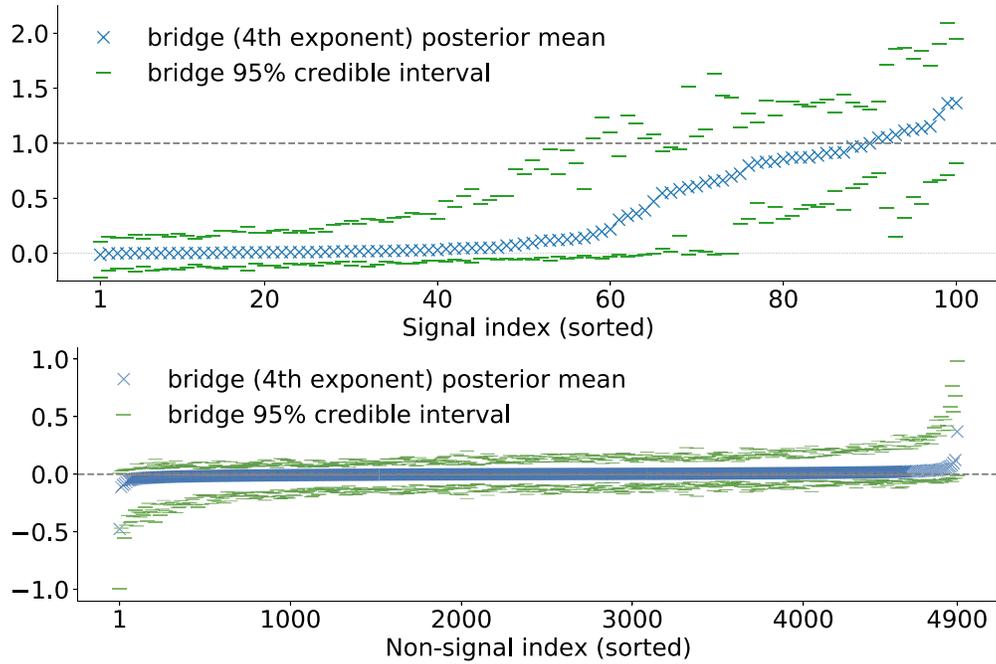


Figure 4.6: The 95% posterior credible intervals under the Bayesian bridge logistic regression with the bridge exponent $1/4$. Compared with the $1/16$ exponent case (Figure 4.4), the posterior distributions have similar means but much wider credible intervals.

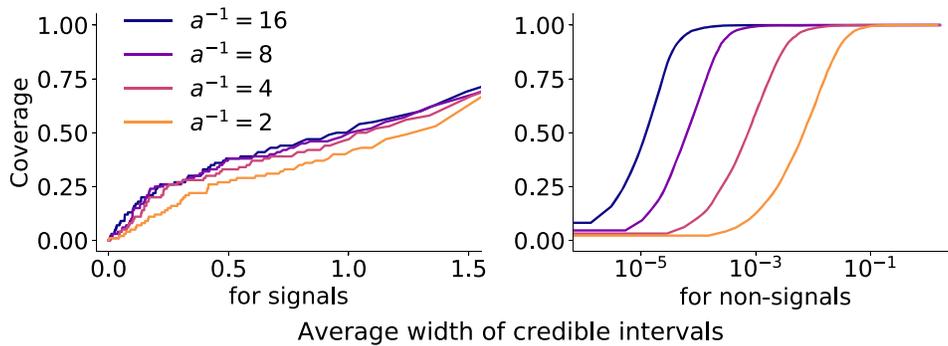


Figure 4.7: Average width v.s. coverage of the credible intervals. The plots are produced by computing the equal-tailed credible intervals at a range of credible levels. The x -axis is in the \log_{10} scale for the non-signals.

intervals are centered around the values similar to the $a = 1/16$ case (Figure 4.4), but are much wider overall. We observe the same pattern throughout the range of the exponent values: similar median values, but tighter intervals for the smaller exponents. In particular, as can be seen in Figure 4.7, more “extreme” shrinkage priors with larger spikes and heavier-tails seem to yield tighter credible intervals for the same coverage.

5 Discussion

Shrinkage priors have been adopted in a variety of Bayesian models, but the potential issues arising from their heavy-tails are often overlooked. Our method provides a simple and convenient way to regularize shrinkage priors, making the posterior inference more robust. Both the theoretical and empirical results demonstrate the benefits of regularization in improving the statistical and computational properties when parameters are only weakly identified. Much of the systematic investigations into the shrinkage prior properties has so far focused on rather simple models and situations in which signals are reasonable strong. Our work adds to the emerging efforts to better understand the behavior of shrinkage models in more complex settings.

Supplementary Material

Supplementary Material of “Shrinkage with Shrunken Shoulders: Gibbs Sampling Shrinkage Model Posteriors with Guaranteed Convergence Rates”
(DOI: [10.1214/22-BA1308SUPP](https://doi.org/10.1214/22-BA1308SUPP); .pdf).

References

- Berger, J. O., Bernardo, J. M., and Sun, D. (2015). “Overall objective priors.” *Bayesian Analysis*, 10(1): 189–221. [MR3420902](https://doi.org/10.1214/14-BA915). doi: <https://doi.org/10.1214/14-BA915>. 381
- Bhadra, A., Datta, J., Polson, N. G., and Willard, B. T. (2017). “Lasso meets horseshoe: a survey.” *Statistical Science*, 34(3): 405–427. [MR4017521](https://doi.org/10.1214/19-STS700). doi: <https://doi.org/10.1214/19-STS700>. 372
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). “Handling sparsity via the horseshoe.” In *Artificial Intelligence and Statistics*, 73–80. 368
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, 97(2): 465–480. [MR2650751](https://doi.org/10.1093/biomet/asq017). doi: <https://doi.org/10.1093/biomet/asq017>. 368
- Choi, H. M. and Hobert, J. P. (2013). “The Polya-Gamma Gibbs sampler for Bayesian logistic regression is uniformly ergodic.” *Electronic Journal of Statistics*, 7: 2054–2064. [MR3091616](https://doi.org/10.1214/13-EJS837). doi: <https://doi.org/10.1214/13-EJS837>. 373, 376
- Flegal, J. M. and Jones, G. L. (2011). “Implementing MCMC: estimating with confidence.” In Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (eds.), *Handbook of Markov chain Monte Carlo*, 175–197. CRC Press. [MR2858449](https://doi.org/10.1214/11-HMC11). 373

- Ghosh, J., Li, Y., and Mitra, R. (2018). “On the use of Cauchy prior distributions for Bayesian logistic regression.” *Bayesian Analysis*, 13(2): 359–383. MR3780427. doi: <https://doi.org/10.1214/17-BA1051>. 368
- Ghosh, P. and Chakrabarti, A. (2017). “Asymptotic optimality of one-group shrinkage priors in sparse high-dimensional problems.” *Bayesian Analysis*, 12(4): 1133–1161. MR3724981. doi: <https://doi.org/10.1214/16-BA1029>. 368
- Greenland, S., Mansournia, M. A., and Altman, D. G. (2016). “Sparse data bias: a problem hiding in plain sight.” *bmj*, 352. 381
- Griffin, J. E. and Brown, P. J. (2010). “Inference with normal-gamma prior distributions in regression problems.” *Bayesian Analysis*, 5(1): 171–188. MR2596440. doi: <https://doi.org/10.1214/10-BA507>. 369
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer. MR2722294. doi: <https://doi.org/10.1007/978-0-387-84858-7>. 385
- Johndrow, J. E., Orenstein, P., and Bhattacharya, A. (2020). “Scalable approximate MCMC algorithms for the horseshoe prior.” *Journal of Machine Learning Research*, 21(73): 1–61. MR4095352. 369, 370, 374, 375, 380
- Jones, G. L. and Hobert, J. P. (2001). “Honest exploration of intractable probability distributions via Markov chain Monte Carlo.” *Statistical Science*, 312–334. MR1888447. doi: <https://doi.org/10.1214/ss/1015346317>. 373
- Kastner, G. (2019). “Sparse Bayesian time-varying covariance estimation in many dimensions.” *Journal of Econometrics*, 210(1): 98–115. MR3944765. doi: <https://doi.org/10.1016/j.jeconom.2018.11.007>. 368
- Kowal, D. R., Matteson, D. S., and Ruppert, D. (2019). “Dynamic shrinkage processes.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(4): 781–804. MR3997101. doi: <https://doi.org/10.1111/rssb.12325>. 368
- Li, K.-C. (1989). “Honest confidence regions for nonparametric regression.” *The Annals of Statistics*, 17(3): 1001–1008. MR1015135. doi: <https://doi.org/10.1214/aos/1176347253>. 384
- Li, Y., Craig, B. A., and Bhadra, A. (2019). “The graphical horseshoe estimator for inverse covariance matrices.” *Journal of Computational and Graphical Statistics*, 28(3): 747–757. MR4007755. doi: <https://doi.org/10.1080/10618600.2019.1575744>. 368
- Louizos, C., Ullrich, K., and Welling, M. (2017). “Bayesian compression for deep learning.” In *Advances in neural information processing systems*, 3288–3298. 368
- Meyn, S. and Tweedie, R. L. (2009). *Markov Chains and Stochastic Stability*. New York, NY, USA: Cambridge University Press. MR2509253. doi: <https://doi.org/10.1017/CBO9780511626630>. 373
- Nishimura, A. and Suchard, M. A. (2022). “Prior-preconditioned conjugate gradient for accelerated Gibbs sampling in ‘large n & large p ’ sparse Bayesian regression.” *Journal*

- of the *American Statistical Association*, advance online publication. doi: <https://doi.org/10.1080/01621459.2022.2057859>. 369, 370, 372
- Nishimura, A. and Suchard, M. A. (2022). “Supplementary Material for ‘Shrinkage with Shrunken Shoulders: Gibbs Sampling Shrinkage Model Posteriors with Guaranteed Convergence Rates.’” *Bayesian Analysis*. doi: <https://doi.org/10.1214/22-BA1308SUPP>. 369, 372, 375, 376, 377, 380, 381
- Pal, S. and Khare, K. (2014). “Geometric ergodicity for Bayesian shrinkage models.” *Electronic Journal of Statistics*, 8(1): 604–645. MR3211026. doi: <https://doi.org/10.1214/14-EJS896>. 369, 374, 375, 376, 380
- Park, T. and Casella, G. (2008). “The Bayesian lasso.” *Journal of the American Statistical Association*, 103(482): 681–686. MR2524001. doi: <https://doi.org/10.1198/016214508000000337>. 368, 385
- Piironen, J. and Vehtari, A. (2017). “Sparsity information and regularization in the horseshoe and other shrinkage priors.” *Electronic Journal of Statistics*, 11(2): 5018–5051. MR3738204. doi: <https://doi.org/10.1214/17-EJS1337SI>. 368, 369, 370, 372, 381
- Polson, N. G. and Scott, J. G. (2010). “Shrink globally, act locally: Sparse Bayesian regularization and prediction.” *Bayesian Statistics*, 9: 501–538. MR3204017. doi: <https://doi.org/10.1093/acprof:oso/9780199694587.003.0017>. 368
- Polson, N. G., Scott, J. G., and Windle, J. (2013). “Bayesian inference for logistic models using Pólya–Gamma latent variables.” *Journal of the American Statistical Association*, 108(504): 1339–1349. MR3174712. doi: <https://doi.org/10.1080/01621459.2013.829001>. 369, 372, 373
- Polson, N. G., Scott, J. G., and Windle, J. (2014). “The Bayesian bridge.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4): 713–733. MR3248673. doi: <https://doi.org/10.1111/rssb.12042>. 368, 369, 377
- Roberts, G. O. and Rosenthal, J. S. (2001). “Markov chains and de-initializing processes.” *Scandinavian Journal of Statistics*, 28(3): 489–504. MR1858413. doi: <https://doi.org/10.1111/1467-9469.00250>. 374
- Roberts, G. O. and Rosenthal, J. S. (2004). “General state space Markov chains and MCMC algorithms.” *Probability Surveys*, 1: 20–71. MR2095565. doi: <https://doi.org/10.1214/154957804100000024>. 373
- Rosenthal, J. S. (1995). “Minorization conditions and convergence rates for Markov chain Monte Carlo.” *Journal of the American Statistical Association*, 90(430): 558–566. MR1340509. 374
- Schuemie, M. J., Ryan, P. B., Hripcsak, G., Madigan, D., and Suchard, M. A. (2018). “Improving reproducibility by using high-throughput observational studies with empirical calibration.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128): 20170356. 370
- Tian, Y., Schuemie, M. J., and Suchard, M. A. (2018). “Evaluating large-scale propen-

- sity score performance through real-world and synthetic data experiments.” *International Journal of Epidemiology*. 381
- van der Pas, S., Salomond, J.-B., and Schmidt-Hieber, J. (2016). “Conditions for posterior contraction in the sparse normal means problem.” *Electronic journal of statistics*, 10(1): 976–1000. MR3486423. doi: <https://doi.org/10.1214/16-EJS1130>. 368
- van der Pas, S., Szabó, B., and van der Vaart, A. (2017). “Adaptive posterior contraction rates for the horseshoe.” *Electronic Journal of Statistics*, 11(2): 3196–3225. MR3705450. doi: <https://doi.org/10.1214/17-EJS1316>. 377, 385
- Wang, X. and Roy, V. (2018). “Geometric ergodicity of Pólya-Gamma Gibbs sampler for Bayesian logistic regression with a flat prior.” *Electronic Journal of Statistics*, 12(2): 3295–3311. MR3861283. doi: <https://doi.org/10.1214/18-EJS1481>. 373, 379

Acknowledgments

We are indebted to Andrew Holbrook for the alliteration in the article title.