

# LAPLACE APPROXIMATION FOR BAYESIAN VARIABLE SELECTION VIA LE CAM'S ONE-STEP PROCEDURE

TIANRUI HOU, LIWEI WANG, AND YVES ATCHADÉ

(May 2024)

ABSTRACT. Variable selection in high-dimensional spaces is a pervasive challenge in contemporary scientific exploration and decision-making. However, existing approaches that are known to enjoy strong statistical guarantees often struggle to cope with the computational demands arising from the high dimensionality. To address this issue, we propose a novel Laplace approximation method based on Le Cam's one-step procedure (OLAP), designed to effectively tackle the computational burden. Under some classical high-dimensional assumptions we show that OLAP is a statistically consistent variable selection procedure. Furthermore, we show that the approach produces a posterior distribution that can be explored in polynomial time using a simple Gibbs sampling algorithm. Toward that polynomial complexity result, we also made some general, noteworthy contributions to the mixing time analysis of Markov chains. We illustrate the method using logistic and Poisson regression models applied to simulated and real data examples.

## 1. INTRODUCTION

Variable selection and model selection methods are crucial tools in the advancement on many data-driven scientific problems (O'Hara and Sillanpää (2009); Fan and Lv (2010)). However existing methods that are known to be statistically consistent in high-dimensional regimes are computationally expensive to deploy (Chen and Chen (2008); Luo and Chen (2011); Barber et al. (2016)). As a result these methods are often implemented using greedy searches that may lack convergence guarantees. We develop in this work a new Laplace approximation that yields a variable selection procedure that is both consistent, and computationally inexpensive to implement. Although our methodology applies more generally, we focus the presentation on a class of generalized linear regression (GLM) models with known dispersion

---

2010 *Mathematics Subject Classification.* 62F15, 62Jxx.

*Key words and phrases.* Regression models, Variable selection, High-dimensional Bayesian inference, Laplace approximation, Markov Chain Monte Carlo mixing times.

This work is partially supported by the NSF grant DMS 2210664.

Tianrui Hou: Boston University, 675 Commonwealth Av., Boston 02215 MA, United States. *E-mail address:* tianruih@bu.edu.

Liwei Wang: Boston University, 675 Commonwealth Av., Boston 02215 MA, United States. *E-mail address:* wlwfoo@bu.edu.

Y. Atchadé: Boston University, 675 Commonwealth Av., Boston 02215 MA, United States. *E-mail address:* atchade@bu.edu.

parameter. Specifically, we assume that we have a data set  $\mathcal{D} \stackrel{\text{def}}{=} \{(y_i, \mathbf{x}_i), 1 \leq i \leq n\}$ , with random response  $y_i \in \mathbb{R}$ , and nonrandom explanatory variable  $\mathbf{x}_i \in \mathbb{R}^p$  collected from  $n$  independent units. We consider a GLM model where, up to additive constant that we ignore, the log-likelihood function is given by

$$(1) \quad \ell(\theta; \mathcal{D}) \stackrel{\text{def}}{=} \sum_{i=1}^n y_i \langle \theta, \mathbf{x}_i \rangle - \psi(\langle \theta, \mathbf{x}_i \rangle), \quad \theta \in \mathbb{R}^p,$$

for some known function  $\psi : \mathbb{R} \rightarrow \mathbb{R}$ , and where  $\theta \in \mathbb{R}^p$  is the parameter of interest. We assume throughout that the model is well-specified, with a true value of the parameter  $\theta_\star \in \mathbb{R}^p$  that satisfies

$$\mathbb{E}(y_i - \psi'(\langle \mathbf{x}_i, \theta_\star \rangle)) = 0,$$

where  $\psi'$  denotes the derivative of  $\psi$ . We consider the classical high-dimensional scenario where  $p$  is potentially much larger than the sample size  $n$ , but  $\theta_\star$  is believed to be sparse. The variable selection problem in this context corresponds to finding the support of  $\theta_\star$ . We take a Bayesian approach and introduce an additional variable  $\delta \in \Delta \stackrel{\text{def}}{=} \{0, 1\}^p$  that encodes the support of  $\theta$ , with prior distribution

$$(2) \quad \pi(\delta) \propto \left(\frac{1}{p^u}\right)^{\|\delta\|_0}, \quad \delta \in \Delta,$$

for some user-defined parameter  $u > 0$ . Furthermore in our prior specification, we assume that given  $\delta$  the components of  $\theta$  are independent, and  $\theta_j | \{\delta_j = 1\}$  follows a Gaussian prior  $\mathbf{N}(0, 1)$ , whereas  $\theta_j | \{\delta_j = 0\}$  follows a degenerate point-mass distribution with full mass at 0. The variable selection problem therefore boils down to sampling from (or computing the mode of) the marginal posterior distribution of  $\delta$  given the data  $\mathcal{D}$  given by

$$(3) \quad \Pi(\delta | \mathcal{D}) \propto \left(\frac{1}{p^u} \sqrt{\frac{1}{2\pi}}\right)^{\|\delta\|_0} \int_{\mathbb{R}^{\|\delta\|_0}} \exp\left(-\frac{1}{2}\|w\|_2^2 + \ell^\delta(w; \mathcal{D})\right) dw,$$

$$\text{where } \ell^\delta(w; \mathcal{D}) \stackrel{\text{def}}{=} \ell((w, 0)_\delta; \mathcal{D}),$$

and where the notation  $(w, 0)_\delta$  denotes the vector of  $\mathbb{R}^p$  obtained from  $w$  by adding 0 at components where  $\delta_j = 0$  (see Section 1.3 below for a more precise definition). When the log-likelihood function  $\ell$  is quadratic, the integrals in (3) have an explicit form, and it is possible to sample from (3) using off-the-shelves MCMC methods on  $\Delta$ . This approach – sometimes with slightly different priors – has been a cornerstone of Bayesian variable selection for several decades (George and McCulloch (1993); O’Hara and Sillanpää (2009); Guan and Stephens (2011); Yang et al. (2016)). Further MCMC advancements to deal with this type of discrete posterior distributions have appeared recently in the literature, although not considered here (Zanella and Roberts (2019); Griffin et al. (2020); Chang and Zhou (2024)).

It is worth noting that this marginalization strategy is generally not applicable beyond the quadratic case. One common solution is data augmentation that samples jointly a model  $\delta$  together with its corresponding parameter, using trans-dimensional MCMC algorithms (Green (2003); Lamnisos et al. (2009)). However, this strategy is known to scale poorly, particularly

because trans-dimensional MCMC algorithms are difficult to tune (Brooks et al. (2003)) and their sampling complexity is poorly understood. A related idea developed in (Carlin and Chib (1995); Atchade and Bhattacharyya (2019)), consists in complementing the parameter set of each model  $\delta$  with additional variables drawn from the so-called pseudo-prior distribution. In the same vein, several carefully constructed priors that are absolutely continuous with respect to the Lebesgue measure have also been proposed and studied in the literature (Polson and Scott (2011); Narisetty and He (2014); Rockova and George (2018); see Bhadra et al. (2019) for a review). These different approaches circumvent the trans-dimensional nature of the posterior distribution, but requires additional tuning of the prior distribution. Furthermore, an important benefit of the marginalization strategy in (3) is lost: integrating out  $\theta$  and sampling only from  $\delta$  (the so-called collapsed Gibbs sampler) typically improves MCMC mixing (see e.g Liu (1994)).

A popular solution in cases where the integrals in (3) are not available in closed form is the use of Laplace approximation. The modes of the resulting posterior approximation are known to be equivalent to the BIC or e-BIC (Schwarz (1978); Chen and Chen (2008)). In (Barber et al. (2016)) it is shown that under some regularity conditions, and for a class of generalized linear models, replacing the integrals in (3) by their Laplace approximations still produces consistent variable selection. Specifically the method replaces the function  $\ell^\delta(\cdot; \mathcal{D})$  by its second order Taylor expansion around the estimator

$$(4) \quad \hat{\theta}^\delta \stackrel{\text{def}}{=} \underset{w \in \mathbb{R}^{\|\delta\|_0}}{\text{Argmax}} \left[ -\frac{1}{2} \|w\|_2^2 + \ell^\delta(w; \mathcal{D}) \right].$$

The major limitation of this approach is that the Laplace approximation requires the computation of the estimators  $\hat{\theta}^\delta$  (as well as the Hessian matrix of  $\ell^\delta$  at  $\hat{\theta}^\delta$ ) for each  $\delta$ , which is typically computationally expensive. In particular in the GLM considered in this work, these estimators are not available in closed form and requires the use of a numerical solver.

In the frequentist realm, the high-dimensional variable selection problem in a GLM model with likelihood such as (1) is often framed as the best-subset selection problem that solves the minimization problem

$$\min -\ell(\theta; \mathcal{D}) \quad \text{subject to} \quad \|\theta\|_0 \leq s,$$

or some relaxation thereof, for a specified sparsity level  $s$  (Beck and Eldar (2013); Hastie et al. (2015)). These estimators have seen renewed interest due to recent optimization advances to tackle an otherwise NP-hard problem (Bertsimas et al. (2016); Hazimeh and Mazumder (2020)). One drawback of this approach is the limited understanding of the statistical properties of the resulting estimators and the lack of effective and adaptive methods for selecting the sparsity level  $s$ .

**1.1. Main contributions.** In this work we revisit the Laplace approximation method outlined above. We circumvent the computational challenge of the method by introducing a simple variant of the Laplace approximation based on Le Cam’s one-step procedure. The one-step procedure that goes back to (Lecam (1969)), is a well-known scheme to improve on

a given estimator through a one step Newton-Raphson update. We refer the reader to (Vaart (1998) Section 5.7) for some basic properties. The method has also seen a resurgence in popularity in recent years in the high-dimensional literature as a way to de-bias high-dimensional estimators (van de Geer et al. (2014); Javanmard and Montanari (2014); Xia et al. (2020)). Hence, starting from an initial estimator  $\tilde{\theta}$  that is computed once, we propose fast approximations of the estimators  $\hat{\theta}^\delta$  for any  $\delta$ , using Le Cam’s one-step procedure. This leads to a new Laplace approximation that we call one-step Laplace approximation (OLAP). Under some basic high-dimension regularity conditions (basically restricted strong convexity, and limited coherence of Hessian matrices) we show that the method has a statistical performance comparable to the standard Laplace approximation with knowledge of the estimators  $\hat{\theta}^\delta$ , while being an order of magnitude faster. We also found empirically that OLAP is statistically more accurate than other state-of-the-art methods such as Skinny-Gibbs (Narisetty et al. (2018)) and SparseVB (Ray et al. (2020)) for variable selection in logistic regression.

We also analyze a simple Gibbs sampler to sample from the posterior distribution of OLAP. Under the same regularity assumptions mentioned above, we show that a plain Gibbs sampler applied to the posterior distribution of OLAP has a mixing time that scales polynomially with  $(n, p)$ . Toward the proof we also derive some general results on the mixing times of Markov Chain Monte Carlo methods that may be of independent interest. In particular we establish in Theorem 8 a new generalization of Cheeger’s inequality for Markov chains that connects the  $\epsilon$ -conductance of Lovász and Simonovits (1993) and the approximate spectral gap of Atchadé (2021).

**1.2. Outline.** We end this introduction with some general notations that are employed throughout the paper. The one-step Laplace approximation of (3) and some basic statistical and computational properties are derived in Section 2. Numerical illustrations of the methods are presented in Section 4, and all the technical proofs are collected in the appendix. We close the paper with a brief summary in Section 5. Section 3 contains some new results on the mixing times of MCMC algorithms that we use in the proof of Theorem 4.

**1.3. Notations.** We also collect here our notations on sparse models. Throughout our parameter space is  $\mathbb{R}^p$  equipped with its Euclidean norm  $\|\cdot\|_2$  and inner product  $\langle \cdot, \cdot \rangle$ . We also use  $\|\cdot\|_0$  which counts the number of non-zero elements, and  $\|\cdot\|_\infty$  which returns the largest absolute value.

We set  $\Delta \stackrel{\text{def}}{=} \{0, 1\}^p$ . Elements of  $\Delta$  are called sparsity structures, or supports. For  $\delta, \delta' \in \Delta$ , we define  $\min(\delta, \delta')$  as the component-wise minimum of the vectors  $\delta$  and  $\delta'$ . And we write  $\delta \subseteq \delta'$  if  $\min(\delta, \delta') = \delta$ . Equivalently,  $\delta \subseteq \delta'$  if  $\delta_j \leq \delta'_j$  for all  $1 \leq j \leq p$ , and we write  $\delta \supseteq \delta'$  if  $\delta' \subseteq \delta$ .

Given  $\delta \in \Delta$ , and  $\theta \in \mathbb{R}^p$ , we write  $\theta_\delta$  to denote the component-wise product of  $\theta$  and  $\delta$ , and  $\delta^c \stackrel{\text{def}}{=} 1 - \delta$ . Assuming  $\|\delta\|_0 = s$ , we will use the notation  $[\theta]_\delta$  to denote the vector  $(\theta_{j_1}, \dots, \theta_{j_s})$ , where  $j_i$  is the  $i$ -th component of  $\delta$  that is non-zero. Conversely, for  $w \in \mathbb{R}^s$ , we define  $(w, 0)_\delta$  as the element of  $\mathbb{R}^p$  such that  $[(w, 0)_\delta]_\delta = w$ . At times we will abuse the notation and use  $\theta_\delta$  and  $[\theta]_\delta$  interchangeably.

## 2. THE ONE-STEP LAPLACE APPROXIMATION

With  $\ell^\delta$  as in (3), we define

$$\bar{\ell}^\delta(w; \mathcal{D}) \stackrel{\text{def}}{=} \ell^\delta(w; \mathcal{D}) - \frac{1}{2} \|w\|_2^2, \quad w \in \mathbb{R}^{\|\delta\|_0}.$$

The Laplace approximation of the posterior distribution  $\Pi(\cdot|\mathcal{D})$  is formed by replacing  $\bar{\ell}^\delta$  in the posterior (3) by its second order Taylor approximation around the mode  $\hat{\theta}^\delta$ , where  $\hat{\theta}^\delta$  is as defined in (4). The resulting Gaussian integral has an explicit form, which leads to the approximation of (3) given by

$$(5) \quad \hat{\Pi}(\delta|\mathcal{D}) \propto \left(\frac{1}{p^\mu}\right)^{\|\delta\|_0} \frac{e^{\bar{\ell}^\delta(\hat{\theta}^\delta; \mathcal{D})}}{\sqrt{\det(\hat{\mathcal{H}}^\delta)}}, \quad \text{where } \hat{\mathcal{H}}^\delta \stackrel{\text{def}}{=} -\nabla^{(2)} \bar{\ell}^\delta(w; \mathcal{D})|_{w=\hat{\theta}^\delta}.$$

Barber et al. (2016) studied the statistical properties of  $\hat{\Pi}$  and showed that it is consistent (in the sense that it puts high probability on the true model) under some regularity conditions, as  $n, p \rightarrow \infty$ . However as observed above, in general the estimator  $\hat{\theta}^\delta$  is computationally expensive to obtain, making the approach difficult to use in large-scale applications.

We propose OLAP, a simple modification based on Le Cam's one-step device. The basic idea is as follows. For any  $v_0 \in \mathbb{R}^{\|\delta\|_0}$ , using a second order Taylor approximation of the function  $\bar{\ell}^\delta(\cdot; \mathcal{D})$  around  $v_0$ , we can approximate the integral  $\int_{\mathbb{R}^{\|\delta\|_0}} e^{\bar{\ell}^\delta(w; \mathcal{D})} dw$  by

$$(6) \quad \frac{(2\pi)^{\frac{\|\delta\|_0}{2}}}{\sqrt{\det(\mathcal{H}^{v_0})}} \exp\left(\bar{\ell}(v_0; \mathcal{D}) + \frac{1}{2}(\mathcal{G}^{v_0})^\top (\mathcal{H}^{v_0})^{-1} \mathcal{G}^{v_0}\right).$$

In the above expression,  $\mathcal{G}^{v_0} \stackrel{\text{def}}{=} \nabla \bar{\ell}^\delta(w; \mathcal{D})|_{w=v_0}$ , and  $\mathcal{H}^{v_0} \stackrel{\text{def}}{=} -\nabla^{(2)} \bar{\ell}^\delta(w; \mathcal{D})|_{w=v_0}$ . The approximation performs at its best when  $v_0$  is near  $\hat{\theta}^\delta$ , the mode of  $\bar{\ell}^\delta$ , and this is what the standard Laplace approximation does. Let  $\tilde{\theta} \in \mathbb{R}^p$  be some initial estimator of  $\theta_\star$ . The estimator  $\tilde{\theta}$  is computed only once, and for any  $\delta \in \Delta$ , we can take  $\tilde{\theta}^\delta \stackrel{\text{def}}{=} [\tilde{\theta}]_\delta$  (where  $[\tilde{\theta}]_\delta$  is the sub-vector of  $\tilde{\theta}$  that collects the components of  $\tilde{\theta}$  for which  $\delta_j = 1$ ) as an initial approximation of  $\hat{\theta}^\delta$ , that we then improve upon using the one-step procedure. Hence our proposed approximation of  $\hat{\theta}^\delta$  is

$$(7) \quad \check{\theta}^\delta \stackrel{\text{def}}{=} \tilde{\theta}^\delta + (\tilde{\mathcal{H}}^\delta)^{-1} \tilde{\mathcal{G}}^\delta,$$

where

$$(8) \quad \tilde{\mathcal{G}}^\delta \stackrel{\text{def}}{=} \nabla \bar{\ell}^\delta(w; \mathcal{D})|_{w=\tilde{\theta}^\delta}, \quad \text{and } \tilde{\mathcal{H}}^\delta \stackrel{\text{def}}{=} -\nabla^{(2)} \bar{\ell}^\delta(w; \mathcal{D})|_{w=\tilde{\theta}^\delta}.$$

The estimator  $\check{\theta}^\delta$  is obtained from  $\tilde{\theta}^\delta$  by a single Newton-Raphson update. In statistical estimation theory, it is well-known that this single step is enough to improve any rate-optimal initial estimator into an efficient rate optimal estimator, by moving it toward the related M-estimator (Vaart (1998); Brouste et al. (2021)). Versions of the procedure have also appeared in the recent statistics literature to de-bias and otherwise improve upon classical high-dimensional estimators (van de Geer et al. (2014); Javanmard and Montanari (2014); Xia et al. (2020)). Hence we expect similar improvements in our context.

Using  $\check{\theta}^\delta$  and (6) leads to the approximation of  $\Pi$  given by

$$(9) \quad \delta \mapsto \left(\frac{1}{p^u}\right)^{\|\delta\|_0} \frac{\exp(\bar{\ell}^\delta(\check{\theta}^\delta; \mathcal{D}) + \frac{1}{2}(\check{\mathcal{G}}^\delta)^\top(\check{\mathcal{H}}^\delta)^{-1}\check{\mathcal{G}}^\delta)}{\sqrt{\det(\check{\mathcal{H}}^\delta)}}, \quad \delta \in \Delta,$$

where  $\check{\mathcal{G}}^\delta$  and  $\check{\mathcal{H}}^\delta$  are defined as in (8), but replacing  $\tilde{\theta}^\delta$  by  $\check{\theta}^\delta$ . We note that if instead of a single update, we iterate (7) a large number of times, then naturally  $\check{\theta}^\delta \approx \hat{\theta}^\delta$ , and  $\check{\mathcal{G}}^\delta \approx 0$ , and (9) becomes (5). Hence we propose to drop the term  $(\check{\mathcal{G}}^\delta)^\top(\check{\mathcal{H}}^\delta)^{-1}\check{\mathcal{G}}^\delta$ , as well as the log-determinant term in (9), leading to OLAP, our proposed Laplace approximation

$$(10) \quad \check{\Pi}(\delta|\mathcal{D}) \propto \left(\frac{1}{p^u}\right)^{\|\delta\|_0} e^{\bar{\ell}^\delta(\check{\theta}^\delta; \mathcal{D})}, \quad \delta \in \Delta.$$

We then solve the variable selection problem by sampling from the OLAP distribution (10).

2.0.1. *Initial estimator  $\tilde{\theta}$ .* Throughout we take the initial estimator  $\tilde{\theta} \in \mathbb{R}^p$  as either lasso, the ridge regression estimator, or more generally the elastic-net estimator defined as

$$\underset{\theta \in \mathbb{R}^p}{\text{Argmin}} \left[ \sum_{i=1}^n -y_i \langle \theta, \mathbf{x}_i \rangle + \psi(\langle \theta, \mathbf{x}_i \rangle) + \lambda_1 \|\theta\|_1 + \frac{\lambda_2}{2} \|\theta\|_2^2 \right],$$

for some well-tuned regularization parameters  $\lambda_1 \geq 0$ ,  $\lambda_2 \geq 0$ . We refer the reader to Hastie et al. (2015) and the extensive literature on these high-dimensional estimators.

2.1. **Statistical properties.** We recall that we have assumed that our data  $\mathcal{D} = \{(y_i, \mathbf{x}_i), 1 \leq i \leq n\}$  is generated from a well-specified GLM model, with true value of the parameter  $\theta_\star$  with support  $\delta_\star$ , and we set  $s_\star \stackrel{\text{def}}{=} \|\theta_\star\|_0$ . In this section we search for conditions under which the OLAP posterior distribution introduced in (10) concentrates around  $\delta_\star$ . We start with the following general definition. Let  $\{\hat{\beta}^\delta, \delta \in \Delta\}$  be some arbitrary family of estimators, where for  $\delta \in \Delta$ ,  $\hat{\beta}^\delta \in \mathbb{R}^{\|\delta\|_0}$ . We say that the family  $\{\hat{\beta}^\delta, \delta \in \Delta\}$  is variable selection consistent if there exists positive constants  $c_1, c_2 < \infty$  such that the following two properties holds.

(1) For all  $\delta \in \Delta$ , and  $\delta_0 \subseteq \delta$ , with  $\min(\delta, \delta_\star) = \delta_0$ , it holds

$$(11) \quad \bar{\ell}(\hat{\beta}^\delta; \mathcal{D}) \leq \bar{\ell}(\hat{\beta}^{\delta_0}; \mathcal{D}) + c_1(\|\delta\|_0 - \|\delta_0\|_0) \log(p),$$

(2) and for all  $\delta_0, \delta \in \Delta$  such that  $\delta_0 \subseteq \delta \subseteq \delta_\star$ , it holds

$$(12) \quad \bar{\ell}(\hat{\beta}^\delta; \mathcal{D}) \geq \bar{\ell}(\hat{\beta}^{\delta_0}; \mathcal{D}) + c_2(\|\delta\|_0 - \|\delta_0\|_0)n.$$

**Remark 1.** This definition basically says that  $\{\hat{\beta}^\delta, \delta \in \Delta\}$  is model selection consistent if increasing a model  $\delta_0$  by adding a non-relevant variable increases the log-likelihood  $\bar{\ell}(\hat{\beta}^{\delta_0}; \mathcal{D})$  only by a factor at most  $\log(p)$ , whereas adding a relevant variable increases the log-likelihood by a factor  $n$ . The behavior depicted in this definition is generally the expected behavior of the (maximum) log-likelihood ratio. However establishing that this behavior prevails for a given model can sometimes be very challenging, particularly in the high-dimensional setting.

Our first result shows that if the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  is variable selection consistent as defined above, then the OLAP distribution (10) puts most of its probability mass on  $\delta_\star$ . For  $j \geq 0$ , we define

$$\mathcal{A}_j \stackrel{\text{def}}{=} \{\delta \in \Delta : \delta \supseteq \delta_\star, \text{ and } \|\delta\|_0 \leq s_\star + j\}.$$

**Theorem 2.** *Assume  $p \geq 2$ . Suppose that the family of one-step estimators  $\{\check{\theta}^\delta, \delta \in \Delta\}$  introduced in (7) is model selection consistent with constants  $c_1, c_2$ . Consider the posterior distribution (10), and suppose that the sparsity parameter  $\mathbf{u}$  satisfies*

$$(13) \quad \mathbf{u} \geq 2(1 + c_1),$$

and the sample size  $n$  satisfies

$$(14) \quad c_2 n \geq 2(\mathbf{u} + 1) \log(p).$$

Then for all  $j \geq 0$ , we have

$$(15) \quad \check{\Pi}(\mathcal{A}_j | \mathcal{D}) \geq 1 - 2 \left( \frac{1}{p^{\frac{\mathbf{u}(j+1)}{2}}} + \frac{2}{e^{c_2 n/4}} \right).$$

*Proof.* See Section A. □

**Remark 3.** We note that the result holds true for  $j = 0$  and yields

$$\check{\Pi}(\delta_\star | \mathcal{D}) \geq 1 - 2 \left( \frac{1}{p^{\frac{\mathbf{u}}{2}}} + \frac{2}{e^{c_2 n/4}} \right).$$

The theorem assumes that the family of one-step estimators  $\{\check{\theta}^\delta, \delta \in \Delta\}$  is model selection consistent. We investigate this key assumption below. An important point to make here is that in fact any family of estimators that is model selection consistent as defined above yields a consistent estimation of  $\delta_\star$ . Specifically, it can be similarly shown that if a family of estimators  $\{\hat{\beta}^\delta, \delta \in \Delta\}$  is model selection consistent then the distribution

$$\delta \mapsto \exp \left( -\mathbf{u} \|\delta\|_0 \log(p) + \bar{\ell}^\delta(\hat{\beta}^\delta; \mathcal{D}) \right), \quad \delta \in \Delta,$$

satisfies the conclusion of the theorem. Of course, this property is of interest only if the estimators  $\hat{\beta}^\delta$  are easy to compute.

**2.2. MCMC sampling and mixing time.** One of the main advantage of OLAP is that the distribution (10) can be easily explored using a simple Gibbs sampler on the discrete space  $\Delta$ . We note that the conditional distribution of  $\delta_j$  given  $\delta_{-j}$  is the Bernoulli distribution  $\mathbf{Ber}(q_j(\delta))$ , with probability of success given by

$$(16) \quad q_j(\delta) \stackrel{\text{def}}{=} \left( 1 + \exp \left( \mathbf{u} \log(p) + \bar{\ell}^{\delta^{(j,0)}}(\check{\theta}_{\delta^{(j,0)}}; \mathcal{D}) - \bar{\ell}^{\delta^{(j,1)}}(\check{\theta}_{\delta^{(j,1)}}; \mathcal{D}) \right) \right)^{-1}.$$

where  $\delta^{(j,0)}$  (resp  $\delta^{(j,1)}$ ) is equal to  $\delta$  except at component  $j$  where  $\delta_j^{(j,0)} = 0$  (resp.  $\delta_j^{(j,1)} = 1$ ). The resulting Gibbs sampler is presented in Algorithm 1.

**Algorithm 1** (Gibbs sampler for OLAP). 

---

Pick  $\delta^{(0)} \in \Delta$  the initial state. Repeat the following steps for  $k = 0, \dots$ . Given  $\delta^{(k)} = \delta \in \Delta$ :

- (1) Set  $\bar{\delta} = \delta$ . Randomly and uniformly select an ordered subset  $J \subset \{1, \dots, p\}$  of size  $J$ .
  - (2) Update  $(\bar{\delta}_{J_1}, \dots, \bar{\delta}_{J_J})$  sequentially: for  $\iota = 1, \dots, J$ , draw  $V_\iota \sim \mathbf{Ber}(q_{J_\iota}(\bar{\delta}))$ , where  $q_j$  is as in (16), and set  $\bar{\delta}_{J_\iota} = V_\iota$ .
  - (3) Set  $\delta^{(k+1)} = \bar{\delta}$ .
- 

The cost of computing  $q_j(\delta)$  is driven by the cost of forming the matrix  $\tilde{\mathcal{H}}^\delta$  in (7), plus the cost of the Cholesky factorization needed to perform the Newton-Raphson update. Hence the computational cost of  $q_j(\delta)$  is of order  $n\|\delta\|_0^2 + \|\delta\|_0^3$ . As a result, we see that the computation cost of the  $k$ -th iteration of Algorithm 1 is of order

$$J \times \max\left(n\|\delta^{(k)}\|_0^2, \|\delta^{(k)}\|_0^3\right).$$

We analyze the mixing time of Algorithm 1. We focus on the case  $J = 1$ , where the resulting Markov chain is reversible and positive. In the case  $J > 1$ , a similar analysis can be developed for the lazy version of the algorithm, but we will not pursue this. Assuming the data  $\mathcal{D}$  fixed, let  $\mathbb{P}(\cdot|\delta)$  denote the probability measure of the Markov chain generated by Algorithm 1 started from  $\delta$ .

**Theorem 4.** *Assume  $p \geq 2$ . Let  $\{\delta^{(k)}, k \geq 0\}$  be the Markov chain generated by Algorithm 1 with  $J = 1$  and initial state  $\delta^{(0)}$ . Suppose that the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  is variable selection consistent with constants  $c_1, c_2$  such that (13) and (14) hold. Suppose also that there exists  $\alpha, 0 \leq \alpha \leq \mathbf{u}(p - s_\star)/4$  such that  $\delta^{(0)}$  satisfies*

$$(17) \quad \check{\mathbb{P}}(\delta^{(0)}|\mathcal{D}) \geq \frac{1}{p^\alpha}.$$

Fix  $\zeta_0 \in (0, 1)$ . Then there exists an absolute constant  $C$  such that if the number of iterations satisfies

$$(18) \quad N \geq C \left( s_\star + \frac{4\alpha}{\mathbf{u}} \right) \left( \log\left(\frac{1}{\zeta_0}\right) + \alpha \log(p) \right) p,$$

then it holds

$$\|\mathbb{P}(\delta^{(N)} = \cdot | \delta^{(0)}) - \check{\mathbb{P}}(\cdot)\|_{\text{tv}} \leq \max\left(\zeta_0, \frac{4}{p^{\frac{\alpha}{4}}}\right).$$

*Proof.* See Section C. □

The main conclusion of the theorem is that under the warm-start condition (17), if the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  is variable selection consistent then Algorithm 1 is fast mixing. Importantly the theorem also highlights the importance of the initial distribution. To provide more insight on this last point, suppose that we slightly strengthen the definition of model selection consistency of  $\{\check{\theta}^\delta, \delta \in \Delta\}$  by replacing (11) and (12) respectively by

$$(19) \quad \bar{\ell}(\check{\theta}^{\delta_0}; \mathcal{D}) \leq \bar{\ell}(\check{\theta}^\delta; \mathcal{D}) \leq \bar{\ell}(\check{\theta}^{\delta_0}; \mathcal{D}) + c_1(\|\delta\|_0 - \|\delta_0\|_0) \log(p),$$



and

$$(20) \quad \bar{\ell}(\check{\theta}^\delta; \mathcal{D}) + c_3(\|\delta_\star\|_0 - \|\delta\|_0)n \geq \bar{\ell}(\check{\theta}_{\delta_\star}; \mathcal{D}) \geq \bar{\ell}(\check{\theta}^\delta; \mathcal{D}) + c_2(\|\delta_\star\|_0 - \|\delta\|_0)n.$$

for some constant  $c_3 \geq c_2$ . Now, suppose that we run Algorithm 1 from some initial state  $\delta^{(0)} \supseteq \delta_\star$ , with  $\|\delta^{(0)}\|_0 = k_0$ . In other words, the initial state has no false-negative. Then, noting that  $\check{\Pi}(\delta_\star | \mathcal{D}) \geq 1/2$  for  $n$  and  $p$  large enough as a consequence of Theorem 2, we have:

$$\check{\Pi}(\delta^{(0)} | \mathcal{D}) = \check{\Pi}(\delta_\star | \mathcal{D}) \frac{\check{\Pi}(\delta^{(0)} | \mathcal{D})}{\check{\Pi}(\delta_\star | \mathcal{D})} \geq \frac{1}{2p^{u(k_0 - s_\star)}} \exp\left(\bar{\ell}(\check{\theta}^{\delta^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}_{\delta_\star}; \mathcal{D})\right).$$

By the first inequality of (19),  $\bar{\ell}(\check{\theta}^{\delta^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}_{\delta_\star}; \mathcal{D}) \geq 0$ . Hence, we conclude that

$$\check{\Pi}(\delta^{(0)} | \mathcal{D}) \geq \frac{1}{2p^{u(k_0 - s_\star)}}.$$

This inequality means that we can apply Theorem 4 with  $\alpha = u(k_0 - s_\star)/2$ , and it follows that Algorithm 1 has a mixing time that is at most  $k_0^2 \log(p) \times p$ . Crucially in this case, the mixing time of the algorithm does not directly depend on the sample size  $n$ .

Consider now the case where the initial state  $\delta^{(0)}$  contains some false negatives. Specifically, suppose that we start Algorithm 1 from  $\delta^{(0)}$  such that  $\|\delta^{(0)}\|_0 = k_0$ , but  $\delta_\star^{(0)} \stackrel{\text{def}}{=} \min(\delta^{(0)}, \delta_\star)$  is a strict sub-model of  $\delta_\star$ , say  $\|\delta_\star^{(0)}\|_0 = s_0 < s_\star$ . In that case, using the definition of  $\check{\Pi}$  in (10), and  $\check{\Pi}(\delta_\star | \mathcal{D}) \geq 1/2$ ,

$$\begin{aligned} \check{\Pi}(\delta^{(0)} | \mathcal{D}) &= \check{\Pi}(\delta_\star | \mathcal{D}) \frac{\check{\Pi}(\delta_\star^{(0)} | \mathcal{D})}{\check{\Pi}(\delta_\star | \mathcal{D})} \frac{\check{\Pi}(\delta^{(0)} | \mathcal{D})}{\check{\Pi}(\delta_\star^{(0)} | \mathcal{D})} \\ &\geq \frac{1}{2p^{u(k_0 - s_\star)}} \exp\left(\bar{\ell}(\check{\theta}^{\delta_\star^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}_{\delta_\star}; \mathcal{D}) + \bar{\ell}(\check{\theta}^{\delta^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}^{\delta_\star^{(0)}}; \mathcal{D})\right). \end{aligned}$$

By the first inequality of (19),  $\bar{\ell}(\check{\theta}^{\delta^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}^{\delta_\star^{(0)}}; \mathcal{D}) \geq 0$ , and the first inequality of (20) yields  $\bar{\ell}(\check{\theta}^{\delta_\star^{(0)}}; \mathcal{D}) - \bar{\ell}(\check{\theta}_{\delta_\star}; \mathcal{D}) \geq -c_3(s_\star - s_0)n$ . Hence it follows that

$$\check{\Pi}(\delta^{(0)} | \mathcal{D}) \geq \frac{\exp(-c_3(s_\star - s_0)n)}{2p^{u(k_0 - s_\star)}} \geq \frac{1}{2p^\alpha},$$

with  $\alpha = u(k_0 - s_\star) + \frac{c_3(s_\star - s_0)n}{\log(p)}$ . In that case, it follows from Theorem 4 that Algorithm 1 has a mixing time that is at most

$$\left(k_0 + \frac{(s_\star - s_0)n}{\log(p)}\right)^2 \log(p) \times p.$$

Here, the mixing time is still polynomial in  $(n, p)$ , however the result suggests that the mixing time is negatively impacted by the sample size  $n$ . This is similar to the worst-case mixing time bound of  $ns_0^2 p \log(p)$  obtained by Yang et al. (2016) for sampling from a version of (5) for variable selection in a high-dimensional linear regression model, where, using their notations,  $s_0$  denotes an upper-bound on  $s_\star$ . Although we have a worst dependence on  $n$ , our work improves on Yang et al. (2016) in two ways. Firstly, instead of a worst-case analysis, our result describes more precisely the effect of the initialization. Secondly, our analysis does not require an a-priori bound on  $s_\star$ . We stress however that our results (as well as Yang

et al. (2016)) only provide upper bounds on the mixing times. It is possible that Algorithm 1 actually converges faster than what is described in Theorem 4.

We illustrate these findings with a Poisson regression simulation example. The data generating process is described in Section 4. Here we set  $p = 1000$  and we vary  $n \in \{1000, 1500, 2000\}$ . The true model has 10 relevant variables. For each  $(n, p)$  we generate 50 datasets, leading to 50 OLAP distributions  $\check{\Pi}(\cdot|\mathcal{D})$ . For each posterior we estimate the mixing time of Algorithm 1 using the  $L$ -lag coupling method of Biswas et al. (2019), employing 30 coupled chains. We compare the mixing times of Algorithm 1 when started from the null model, and a version started from  $\delta^{(0)}$  that contains the true model plus 10 false-positive. The boxplots of the distributions of the mixing times are given in Figure 1. The results are consistent with our theory. This finding highlights the importance of good MCMC initialization (warm-start). In all the simulation below we initialize the algorithm from the support of the lasso estimator, which under mild conditions is known to contain the true support for sufficiently strong signal (Meinshausen and Yu (2009)).

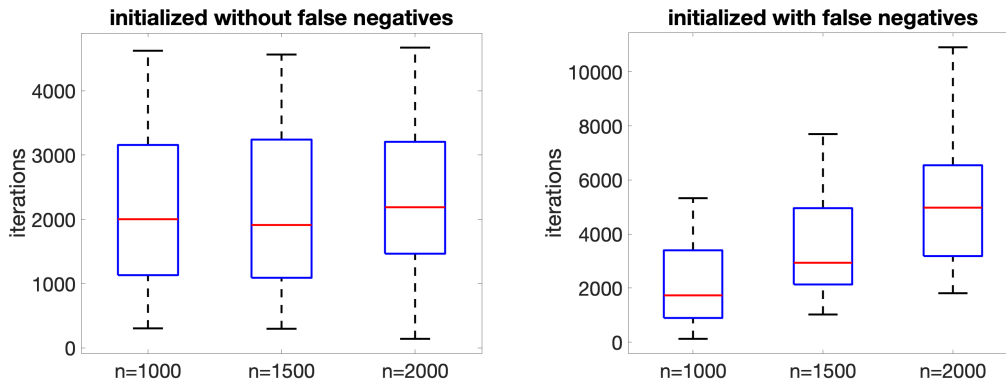


FIGURE 1. Estimated mixing times of Algorithm 1 under different initialization  $\delta^{(0)}$  and for different data size  $n$ . Left:  $\delta^{(0)}$  contains all relevant variables and 10 false-positive. Right:  $\delta^{(0)}$  is null model.

**2.3. Variable selection consistency of the one-step family.** Since the cost per iteration of Algorithm 1 is polynomial in  $(n, p)$ , Theorem 2 and Theorem 4 together show that in problems where the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  is variable selection consistent, the variable selection problem can be provably solved with a computational cost that is polynomial in  $(n, p)$  using OLAP. But under what condition is the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  variable selection consistent? We show in this section that if the initial estimator  $\tilde{\theta}$  is rate optimal (lasso is known to be rate optimal under mild condition), then the one-step family  $\{\check{\theta}^\delta, \delta \in \Delta\}$  inherits the variable selection consistency of the MLE family  $\{\hat{\theta}^\delta, \delta \in \Delta\}$ . This result mirrors the way in which the one-step estimator inherits the efficiency of the MLE in regular models (Vaart (1998) Section 5.7). The variable selection consistency of the MLE family is a more classical problem that has been considered in the literature. For instance, for a class of sparse

GLM models, it is shown to hold in Barber et al. (2016) Theorem 2.2. We work under the following basic set of assumptions on the data generating process.

**H1.** (1) *There exists an absolute constant  $b < \infty$  such that*

$$(21) \quad \max_{1 \leq i \leq n} \max_{1 \leq j \leq p} |\mathbf{x}_{ij}| \leq b.$$

(2) *There exists an absolute constant  $M_0 < \infty$  such that*

$$\left\| \sum_{i=1}^n (y_i - \psi'(\langle \theta_*, \mathbf{x}_i \rangle)) \mathbf{x}_i \right\|_{\infty} \leq M_0 \sqrt{n \log(p)}.$$

(3) *The function  $\psi$  is strictly convex and there exists an absolute constant  $c_3 > 0$  such that for all  $x \in \mathbb{R}$ ,*

$$\left| \frac{d^3 \psi(x)}{dx^3} \right| \leq c_3 \frac{d^2 \psi(x)}{dx^2}.$$

(4) *There exists  $0 < \underline{\kappa} \leq \bar{\kappa}$  such that for all  $\delta_0 \subseteq \delta_*$ , for all  $w \in \mathbb{R}^p$ , with  $\|w\|_0 \leq s_*$  with  $\|w\|_2 = 1$ , it holds*

$$\underline{\kappa} n \leq w^T \left( \sum_{i=1}^n \psi''(\langle \mathbf{x}_i, \tilde{\theta}^{\delta_0} \rangle) \mathbf{x}_i \mathbf{x}_i^T \right) w \leq \bar{\kappa} n.$$

(5) *There exist constants  $c$ , and  $M$  such that for all  $1 \leq j \neq k \leq p$ ,*

$$\sup_{\vartheta \in \mathbb{R}^p: \|\vartheta\|_0 \leq s_*, \|\vartheta - \theta_*\|_2 \leq c \sqrt{\frac{s_* \log(p)}{n}}} \left| \sum_{i=1}^n \psi''(\langle \vartheta, \mathbf{x}_i \rangle) \mathbf{x}_{ij} \mathbf{x}_{ik} \right| \leq M \sqrt{n \log(p)}.$$

**Remark 5.** H1-(2) holds if  $y_i - \psi'(\langle \theta_*, \mathbf{x}_i \rangle)$  is mean-zero and sub-Gaussian, which holds for many GLM models, including linear, logistic, and Poisson regression. Indeed, if data  $\mathcal{D} \stackrel{\text{def}}{=} \{(y_i, \mathbf{x}_i), 1 \leq i \leq n\}$  is a sequence of independent and identically distributed random variables, and the sub-Gaussian norm of  $y_i - \psi'(\langle \theta_*, \mathbf{x}_i \rangle)$  is  $\sigma^2$ , then by Hoeffding's inequality (Vershynin (2018) Theorem 2.6.2), and a union bound inequality, there exists an absolute constant  $c_0$  such that

$$(22) \quad \mathbb{P} \left( \left\| \sum_{i=1}^n (y_i - \psi'(\langle \theta_*, \mathbf{x}_i \rangle)) \mathbf{x}_i \right\|_{\infty} > c_0 \sigma \sqrt{bn \log(p)} \right) \leq \frac{2}{p}.$$

Assumption H1-(3) is a self-concordance assumption that is satisfied by commonly used functions including in linear, logistic and Poisson regression models. For instance for logistic regression,  $\psi(x) = \log(1 + e^x)$ , and

$$\left| \frac{d^3 \psi(x)}{dx^3} \right| = \left| \frac{e^x (e^x - 1)}{(1 + e^x)^3} \right| \leq \frac{e^x}{(1 + e^x)^2} = \frac{d^2 \psi(x)}{dx^2}.$$

Hence H1-(3) holds with  $c_3 = 1$ . For Poisson regression  $\psi(x) = e^x$ , hence H1-(3) also holds with  $c_3 = 1$ .

H1-(4) can be shown to follow from a restricted eigenvalue assumption on the sample covariance matrix  $X^T X/n$ , where  $X \in \mathbb{R}^{n \times p}$  is the matrix of covariates with  $i$ -th row given

by  $\mathbf{x}'_i$ . To see this, note that for most reasonable initial estimators, with high probability it holds  $\|\tilde{\theta}\|_1 \leq b_2$  for some constant  $b_2$  that can be assumed independent of  $(n, p)$ . Therefore in that case, using H1-(1),  $|\langle \mathbf{x}_i, \tilde{\theta}^{\delta_0} \rangle| \leq b\|\tilde{\theta}\|_1 \leq B$  for some constant  $B$ . Thus, letting  $\tau = \inf_{x \in [-B, B]} \psi''(x) > 0$ , we see that with high probability we have

$$w^\top \left( \sum_{i=1}^n \psi''(\langle \mathbf{x}_i, \tilde{\theta}^{\delta_0} \rangle) \mathbf{x}_i \mathbf{x}_i^\top \right) w = \sum_{i=1}^n \psi''(\langle \mathbf{x}_i, \tilde{\theta}^{\delta_0} \rangle) \langle \mathbf{x}_i, w \rangle^2 \geq \tau w^\top (X^\top X) w.$$

The restricted eigenvalue assumption on the sample covariance matrix  $X^\top X/n$  is known to hold in setting where the matrix  $X$  can be viewed as a realization of a random matrix with i.i.d rows drawn from a sub-Gaussian distribution. We refer the reader for instance to Wainwright (2019) Chapter 9 for more details.

H1-(5) imposes a bound on the off-diagonal elements of the Hessian matrices of the log-likelihood. In the linear case, this corresponds to the limited coherence assumption of the design matrix commonly imposed (Candès and Plan (2009)). See also Foygel and Drton (2015) for a similar assumption in a logistic regression setting.

**Theorem 6.** *Assume H1, with  $p \geq 4$ . Suppose that the initial estimator  $\tilde{\theta}$  and the oracle MLE  $\hat{\theta}^{\delta_\star}$  satisfy*

$$(23) \quad \max \left( \|\hat{\theta}^{\delta_\star} - \theta_\star\|_2, \|\tilde{\theta} - \theta_\star\|_2 \right) \leq C \sqrt{\frac{s_\star \log(p)}{n}},$$

for some constant  $C$ . If the MLE family  $\{\hat{\theta}^\delta, \delta \in \Delta\}$  is model selection consistent, then we can find a constant  $C'$  such that with a sample size

$$(24) \quad n \geq C' s_\star^4 \log(p),$$

the one-step family  $\{\hat{\theta}^\delta, \delta \in \Delta\}$  is also model selection consistent.

*Proof.* See Section B. □

### 3. APPROXIMATE SPECTRAL GAPS FOR MARKOV KERNELS

The proof of Theorem 4 is based on some general results of independent interest that we present in this section, and which provide lower bounds on the mixing time of Markov kernels. Specifically, Theorem 4 is proved by applying Proposition 9 and Lemma 7 below to the Gibbs sampler chain produced by Algorithm 1.

Let  $\pi$  be a density on some measure space  $\mathcal{X}$  equipped with a sigma-algebra  $\mathcal{B}$  and a reference sigma-finite measure  $dx$ . We will also write  $\pi$  to denote the induced probability measure:  $\pi(dx) = \pi(x)dx$ . For a function  $f : \mathcal{X} \rightarrow \mathbb{R}$ , we write  $f \in \mathcal{B}$  to say that  $f$  is  $\mathcal{B}$ -measurable. We also define

$$\|f\|_\infty \stackrel{\text{def}}{=} \sup_{x \in \mathcal{X}} |f(x)|, \quad \text{and} \quad \text{osc}(f) \stackrel{\text{def}}{=} \sup_{y, z \in \mathcal{X}} |f(y) - f(z)|.$$

We let  $L^2(\pi)$  denote the Hilbert space of all real-valued square-integrable (wrt  $\pi$ ) functions on  $\mathcal{X}$ , equipped with the inner product  $\langle f, g \rangle \stackrel{\text{def}}{=} \int_{\mathcal{X}} f(x)g(x)\pi(dx)$  with associated norm

$\|\cdot\|_2$ . We will also make use of the essential supremum (resp. essential infimum) of  $f$  with respect to  $\pi$  defined as  $\mathbf{ess-sup}(f) \stackrel{\text{def}}{=} \inf\{M \geq 0 : \pi(\{x \in \mathcal{X} : |f(x)| > M\}) = 0\}$  (resp.  $\mathbf{ess-inf}(f) \stackrel{\text{def}}{=} \sup\{M \geq 0 : \pi(\{x \in \mathcal{X} : |f(x)| < M\}) = 0\}$ ).

If  $K$  is a Markov kernel on  $\mathcal{X}$ , and  $n \geq 1$  an integer,  $K^n$  denotes the  $n$ -th iterate of  $K$ , defined recursively as  $K^n(x, A) \stackrel{\text{def}}{=} \int_{\mathcal{X}} K^{n-1}(x, dz)K(z, A)$ ,  $x \in \mathcal{X}$ ,  $A$  measurable. For  $f \in \mathcal{B}$ , we define  $Kf : \mathcal{X} \rightarrow \mathbb{R}$  as  $Kf(x) \stackrel{\text{def}}{=} \int_{\mathcal{X}} K(x, dz)f(z)$ ,  $x \in \mathcal{X}$ , whenever the integral is well defined. And if  $\mu$  is a probability measure on  $\mathcal{X}$ , then  $\mu K$  is the probability on  $\mathcal{X}$  defined as  $\mu K(A) \stackrel{\text{def}}{=} \int_{\mathcal{X}} \mu(dz)K(z, A)$ ,  $A \in \mathcal{B}$ . The total variation distance between two probability measures  $\mu, \nu$  is defined as

$$\|\mu - \nu\|_{\text{tv}} \stackrel{\text{def}}{=} 2 \sup_{A \in \mathcal{B}} (\mu(A) - \nu(A)) = \sup_{f \in \mathcal{B}: \|f\|_{\infty} \leq 1} \left( \int_{\mathcal{X}} f(x)\mu(dx) - \int_{\mathcal{X}} f(x)\nu(dx) \right).$$

If the Markov kernel  $K$  has invariant distribution  $\pi$ , without changing notation we will also view  $K$  as the linear operator on  $L^2(\pi)$  that transforms  $f$  to  $Kf$  as defined above. We write  $K^*$  to denote the adjoint of  $K$ , that is the linear operator on  $L^2(\pi)$  such that  $\langle Kf, g \rangle = \langle f, K^*g \rangle$  for all  $f, g \in L^2(\pi)$ . We say that  $K$  is reversible with respect to  $\pi$  ( $\pi$ -reversible for short) if  $K^* = K$ , and we say that  $K$  is positive if it is  $\pi$ -reversible and  $\langle f, Kf \rangle \geq 0$  for all  $f \in L^2(\pi)$ . For  $f \in L^2(\pi)$ , we set  $\pi(f) \stackrel{\text{def}}{=} \int_{\mathcal{X}} f(x)\pi(dx)$ ,  $\text{Var}_{\pi}(f) \stackrel{\text{def}}{=} \|f - \pi(f)\|_2^2$ , and

$$\mathcal{E}_K(f, f) \stackrel{\text{def}}{=} \frac{1}{2} \int_{\mathcal{X}} \int_{\mathcal{X}} (f(y) - f(x))^2 \pi(dx)K(x, dy) = \langle f, f \rangle - \langle f, Kf \rangle.$$

To quantify the rate of convergence of  $K^n$  towards  $\pi$ , we will use the concept of approximate spectral gap (Atchadé (2021)). For  $\zeta \in [0, 1)$ , the  $\zeta$ -spectral gap of  $K$  is

$$(25) \quad \lambda_{\zeta}(K) \stackrel{\text{def}}{=} \inf \left\{ \frac{\mathcal{E}_K(f, f)}{\text{Var}_{\pi}(f) - \frac{\zeta}{2}}, f \in L^2(\pi) \text{ s.t. } \|f\|_{\infty} \leq 1 \text{ and } \text{Var}_{\pi}(f) > \zeta \right\}.$$

We recover the classical spectral gap (denoted  $\lambda(K)$ ) by taking  $\zeta = 0$ , and replacing the infinity norm by the  $L^2$ -norm<sup>1</sup>. We note that when  $K$  is positive then  $0 \leq \lambda_{\zeta}(K) \leq 2$ . This quantity can be used to quantify the convergence rate of  $K$  toward a total variation ball of radius  $O(\sqrt{\zeta})$  around  $\pi$ . Specifically, the following is a slight modification of Lemma 2.1 of Atchadé (2021). We provide a proof for completely. The reversibility or positivity of  $K$  are not needed, but are imposed here for simplicity.

**Lemma 7.** *Suppose that  $K$  is  $\pi$ -reversible and positive, and let  $\zeta \in [0, 1)$ . Let  $\pi_0(dx) = f_0(x)\pi(dx)$ , for some bounded function  $f_0$ . For all integer  $N \geq 1$ , we have*

$$(26) \quad \|\pi_0 K^N - \pi\|_{\text{tv}}^2 \leq \max \left( \zeta \|f_0\|_{\infty}^2, \left[ 1 - \frac{\lambda_{\zeta}(K)}{2} \right]^N \text{Var}_{\pi}(f_0) \right).$$

*Proof.* See Section D.1. □

<sup>1</sup>In the definition of  $\lambda_{\zeta}(K)$  one can replace the infinity norm  $\|\cdot\|_{\infty}$  by any other norm that satisfies  $\|Kf\| \leq \|f\|$  for all  $f \in L^2(\pi)$ . We use the infinity norm here mainly for convenience.

For  $N \geq 2 \log(\zeta^{-1})/\lambda_\zeta(K)$ , we have

$$\left(1 - \frac{\lambda_\zeta(K)}{2}\right)^N \leq \exp\left(-\frac{N\lambda_\zeta(K)}{2}\right) \leq \zeta.$$

Hence it follows from the theorem that for  $N \geq 2 \log(\zeta^{-1})/\lambda_\zeta(K)$ , we have  $\|\pi_0 K^N - \pi\|_{\text{tv}} \leq \|f_0\|_\infty \sqrt{\zeta}$ . We note however that the right-hand side of (26) does not converge to 0 as  $N \rightarrow \infty$ : unlike the spectral gap, the approximate spectral gap measures only converge to within  $O(\sqrt{\zeta})$  of  $\pi$ . On the plus side, the method has the advantage that in many problems it scales better with the dimension of the problem than  $\lambda(K)$ .

**3.1. Conductance and Cheeger's inequality.** A similar concept that pre-dates and has motivated the development of the approximate spectral gap is the  $\zeta$ -conductance of Lovász and Simonovits (1993). For  $\zeta \in [0, 1/2)$ , the  $\zeta$ -conductance of the Markov kernel  $K$  is defined as

$$\Phi_\zeta(K) \stackrel{\text{def}}{=} \inf \left\{ \frac{\int_A \pi(\text{d}x) K(x, A^c)}{(\pi(A) - \zeta)(\pi(A^c) - \zeta)}, A \in \mathcal{B} : \zeta < \pi(A) < 1 - \zeta \right\}.$$

The special case  $\Phi_0(K)$  (that is,  $\zeta = 0$ ) corresponds to the standard conductance (Lawler and Sokal (1988); Sinclair and Jerrum (1989); Douc et al. (2018)), and we will also denote it by  $\Phi(K)$ . The conductance  $\Phi(K)$  captures how rapidly a Markov chain in stationarity is able to move around the space.  $\Phi_\zeta(K)$  measures a similar feature but ignore small sets. Corollary 1.5 of Lovász and Simonovits (1993) shows that with a warm start, a Markov chain with transition kernel  $K$  converges to within  $O(\zeta)$  of  $\pi$  in at most  $2 \log(1/\zeta)/\Phi_\zeta(K)^2$  number of iterations. This is analogous to Lemma 7. Here too, the  $\zeta$ -conductance often depends more favorably on the dimensional of the problem than  $\Phi(K)$ , however it only measures convergence to within  $O(\zeta)$  to  $\pi$ . In many problems, particularly when dealing with log-concave distributions where a rich set isoperimetric inequalities are available, the conductances are typically easier to control than the spectral gaps (see for instance Chen et al. (2020) for references and some recent results).

The relationship between the conductance and the spectral gap has been investigated in the literature, and is captured by Cheeger's inequality (Lawler and Sokal (1988)) that states that

$$(27) \quad \frac{\Phi(K)^2}{8} \leq \lambda(K) \leq \Phi(K).$$

A similar relationship is expected to hold between  $\Phi_\zeta(K)$  and  $\lambda_\zeta(K)$ . We close the gap in the literature by showing that this is indeed the case.

**Theorem 8.** *For all  $\epsilon \in [0, 1/2)$ . We have*

$$\frac{(\Phi_{\frac{\epsilon}{32}}(K))^2}{16} \leq \lambda_{\frac{\epsilon}{2}}(K) \leq \Phi_\epsilon(K).$$

*Proof.* See Section D.2. □

**3.2. Bounds using canonical paths.** Moving closer to our intended application, we assume now that  $\mathcal{X}$  is a discrete set. We recall some basic definitions from graph theory. A graph  $(\mathcal{V}, \mathcal{E})$  with vertex set  $\mathcal{V}$  and edge set  $\mathcal{E}$  is a set  $\mathcal{V}$  together with a subset  $\mathcal{E}$  of  $\mathcal{V} \times \mathcal{V}$ . We say that  $(\mathcal{V}, \mathcal{E})$  is undirected if for all  $(x, y) \in \mathcal{V} \times \mathcal{V}$ ,  $(x, y) \in \mathcal{E}$  if and only if  $(y, x) \in \mathcal{E}$ . We will write an edge as  $(e_-, e_+)$ , where  $e_-, e_+$  denote its two incident nodes. We say that the graph is connected if for all  $(x, y) \in \mathcal{V} \times \mathcal{V}$ ,  $x \neq y$ , we can find a sequence of edges  $(z_0, z_1), \dots, (z_{\ell-1}, z_\ell)$  such that  $z_0 = x$ ,  $z_\ell = y$ , and  $(z_{k-1}, z_k) \in \mathcal{E}$  for  $k = 1, \dots, \ell$ . The integer  $\ell$  is the length of the path. There may exist many paths linking any two points. For each pair  $(x, y) \in \mathcal{V}$ ,  $x \neq y$ , we assume given a special path denoted  $\gamma_{xy}$  linking  $(x, y)$  that we call a canonical path. We impose the additional restriction that an edge can appear only once along a given canonical path. We write  $|\gamma_{xy}|$  to denote the length of the canonical path  $\gamma_{xy}$ , and  $\Gamma \stackrel{\text{def}}{=} \{\gamma_{xy}, x, y \in \mathcal{V}\}$  for the set of all canonical paths. We make the following assumption.

**H2.** *There exists a subset  $\mathcal{X}_0 \subseteq \mathcal{X}$  such that  $\pi(\mathcal{X}_0) > 0$ , and a connected undirected graph  $(\mathcal{X}_0, \mathcal{E}_0)$  with canonical paths  $\Gamma \stackrel{\text{def}}{=} \{\gamma_{xy}, x, y \in \mathcal{X}_0\}$  such that for all  $(x, y) \in \mathcal{E}_0$ ,  $\pi(x)K(x, y) > 0$ .*

We define

$$(28) \quad \mathfrak{m}(\mathcal{X}_0) \stackrel{\text{def}}{=} \max_{e \in \mathcal{E}_0} \sum_{\gamma_{xy}: \gamma_{xy} \ni e} \frac{|\gamma_{xy}| \pi(x) \pi(y)}{\pi(e_-) K(e_-, e_+)}.$$

When  $\mathcal{X}_0 = \mathcal{X}$ ,  $\mathfrak{m}(\mathcal{X}_0)$  is the geometric measure of bottleneck of Sinclair (1992) (see also Diaconis and Stroock (1991)). In many cases by carefully choosing  $\mathcal{X}_0$ ,  $\mathfrak{m}(\mathcal{X}_0)$  scales better than  $\mathfrak{m}(\mathcal{X})$ . The next result is analogous to Proposition 1 of Diaconis and Stroock (1991).

**Proposition 9.** *Assume H2. Given  $\epsilon \in [0, 1/2)$ , if  $\pi(\mathcal{X}_0) \geq 1 - \epsilon/8$ , then  $\lambda_\epsilon(K) \geq \mathfrak{m}(\mathcal{X}_0)^{-1}$ .*

*Proof.* See Section D.3. □

**Remark 10.** Theorem 4 is proved by applying Proposition 9 and Lemma 7 to the Gibbs sampler chain produced by Algorithm 1.

#### 4. NUMERICAL ILLUSTRATION

We investigate several aspects of Algorithm 1 in a simulation setting using logistic and Poisson regression. Here is the simulation setup. We generate  $X \in \mathbb{R}^{n \times p}$  with independent rows drawn from  $\mathbf{N}_p(0, \Sigma)$ , where  $\Sigma_{kj} = \varrho^{|j-k|}$ , where  $\varrho \in \{0, 0.9\}$  referred to as “low correlation” and as “highly correlation”, respectively. We consider various combinations of  $n, p$ .

As a true  $\theta_* \in \mathbb{R}^p$ , we use a sparse  $\theta_*$  with first 10 components uniformly drawn from  $(-3, -2) \cup (2, 3)$ . For logistic regression, we draw the response as  $Y_i \sim \mathbf{Ber}(p_i)$ , with  $p_i = (1 + \exp - \langle \mathbf{x}_i, \theta_* \rangle)^{-1}$ , where  $\mathbf{x}_i$  denotes the  $i$ -th row of  $X$ . For Poisson regression we draw the response as  $Y_i \sim \mathbf{Poi}(\lambda_i)$ , where  $\lambda_i = e^{\langle \mathbf{x}_i, \theta_* \rangle}$ .

Throughout we set the prior parameter  $u = 0.8$ , and in Algorithm 1, we set  $J = 100$ , and unless stated otherwise, we take  $\delta^{(0)}$  as the support of **lasso**.

#### 4.1. Illustration using logistic regression.

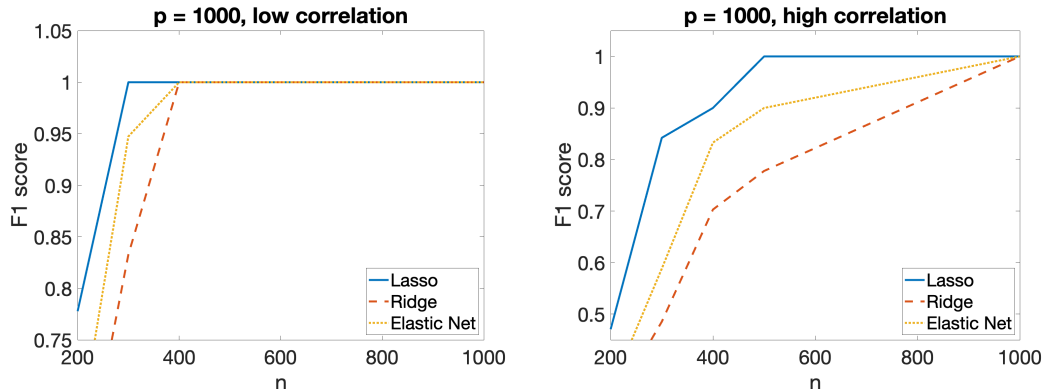


FIGURE 2. Comparison of F1-scores for different initial estimators for logistic models.

4.1.1. *Effect of the initial estimator.* We first investigate the effect of the initial estimator  $\tilde{\theta}$  on OLAP. We compare three different initializations: **lasso**, **ridge**, and **elastic net**. We utilize R package `glmnet` to calculate these initial estimators. As a measure of performance we compute the F1-score (harmonic mean of sensitivity and precision) along the MCMC chain.

We focus on the case when  $p = 1000$  and we increase the sample size  $n$  from 200 to 1000, under both the low and high correlation settings. Under each simulation set up, we run Algorithm 1 50 times, and we report the median of F1-scores of the 50 chains after mixing.

As expected we observe from Figure 2 that the F1 score is an increasing function of the sample size, and in all scenarios, the **lasso** initialization produces the highest F1-scores. We also observe that given enough sample size, all three initializations perform well, which is consistent with our theoretical findings. Based on these results, we focus on the **lasso** initialization for the remaining experiments.

4.1.2. *Mixing time comparisons.* We compare empirically the mixing time of Algorithm 1 with the mixing time of the exact method that employs the data-augmentation strategy proposed in Atchade and Bhattacharyya (2019). The method consists in sampling jointly  $(\delta, \theta)$  from the distribution

$$(29) \quad \Pi(\delta, \theta | \mathcal{D}) \propto \left( \frac{1}{p^u} \sqrt{\frac{1}{\rho_0}} \right)^{\|\delta\|_0} \exp \left( -\frac{1}{2} \|\theta_\delta\|_2^2 - \frac{\rho_0}{2} \|\theta - \theta_\delta\|_2^2 + \ell(\theta_\delta; \mathcal{D}) \right),$$

for some hyper-parameter  $\rho_0$ . One can sample from this distribution by Metropolis-within-Gibbs, alternating between a Gibbs update on  $\delta$  given  $\theta$ , and a Metropolis-Hastings update on  $\theta$  given  $\delta$  (here we use MaLA). We note that the marginal distribution of  $\delta$  under (29) is precisely (3). So the method is exact, and we refer to it as the Exact method. For both algorithms we estimate their mixing times using the  $L$ -lag coupling method proposed by Biswas et al. (2019). The  $L$ -lag coupling method consists in running coupled version of an MCMC algorithm until a coupling event. The mixing time of the algorithm can then be



related to a moment of the coupling time. We refer the reader to Biswas et al. (2019) and Jacob et al. (2020). In our implementation of the  $L$ -lag coupling method we couple the chains using a straightforward maximal coupling of the Bernoulli draws in Algorithm 1. For coupling the Metropolis-within-Gibbs sampling for (29) we follow the algorithm in Atchade and Wang (2023).




FIGURE 3. Iterations and mixing time (in seconds) of logistic regression, low correlation.




FIGURE 4. Iterations and mixing time (in seconds) of logistic regression, high correlation.

For the comparison we set  $p = 1000$  and  $n = 1000$ , and we consider both the low and high correlation cases. In each setting, we generate 50 datasets. Each dataset defines a posterior distribution against which the algorithms under consideration have a mixing time. We estimate these mixing times by running 30 coupled chains. We could observe from the boxplots on Figures 3 and 4 that Algorithm 1 mixes in a much smaller number of iterations and has a shorter running time compared to the Metropolis-within-Gibbs sampler for the Exact method. Specifically, in the low correlation case, OLAP sampler has a median burn-in of 60 iterations, and a median burn-in running time of 1.5 seconds, and our sampler for the Exact method has a median burn-in of 408 iterations, and a median burn-in running time of 4.7 seconds. In the high correlation case, OLAP sampler has a median burn-in of 60.5

iterations, and a median burn-in running time of 1.4 seconds, and the sampler for the Exact method has a median burn-in of 753 iterations, and a median burn-in running time of 8.6 seconds.

TABLE 1. F1-score for logistic regression.  $p = 1000$ , low correlation

n	S-Gibbs		SparseVB		One-step Lasso		Exact		Lasso	
	Median	Std. Error	Median	Std. Error	Median	Std. Error	Median	Std. Error	Median	Std. Error
200	0.778	0.142	0.889	0.092	0.778	0.093	0.750	0.220	0.300	0.080
300	0.909	0.067	1.000	0.009	1.000	0.052	1.000	0.035	0.294	0.083
400	0.909	0.061	1.000	0.025	1.000	0	1.000	0	0.280	0.098
500	0.909	0.065	1.000	0.025	1.000	0	1.000	0	0.288	0.097
1000	0.952	0.058	1.000	0.026	1.000	0	1.000	0	0.339	0.122

4.1.3. *Statistical performance and comparison with other methods.* Here we perform a comparison between OLAP, the Exact method, Skinny-Gibbs (Narisetty et al. (2018)), SparseVB (Ray et al. (2020)), and the standard **lasso** in terms of statistical recovery of the true support  $\delta_*$ . For the simulation we use the R packages `skinnybasad` (for Skinny-Gibbs) and `svb` (for SparseVB) provided by the authors and we use the popular `glmnet` for **lasso**. We focus on the scenarios with  $p = 1000$ , and vary the sample sizes and correlations. We report both the medians and standard errors of the F1-scores over 50 chains. The results are presented in Table 1 and 2.

We could observe from Table 1 that when  $\rho = 0$ , OLAP, the Exact method, and SparseVB have roughly the same F1 score, and are all better than S-Gibbs and LASSO. While, with high correlation data, we could observe from Table 2 that OLAP matches the exact method, and both are better than SparseVB and S-Gibbs.

TABLE 2. F1-score for logistic regression.  $p = 1000$ , high correlation

n	S-Gibbs		SparseVB		One-step Lasso		Exact		Lasso	
	Median	Std. Error	Median	Std. Error	Median	Std. Error	Median	Std. Error	Median	Std. Error
200	0.358	0.169	0.500	0.167	0.471	0.175	0.572	0.184	0.229	0.064
300	0.476	0.181	0.718	0.165	0.842	0.164	0.842	0.136	0.265	0.070
400	0.586	0.162	0.800	0.143	0.900	0.096	0.900	0.105	0.281	0.076
500	0.652	0.172	0.842	0.105	1.000	0.070	0.947	0.071	0.314	0.057
1000	0.714	0.142	1.000	0.082	1.000	0	1.000	0.037	0.310	0.052

4.2. **Numerical illustration using Poisson regression.** Similarly, we investigate several aspects of Algorithm 1 with a simulated Poisson regression example. The setting is very similar to the logistic experiments, except that, we draw the response as  $y_i \sim \mathbf{Poi}(\lambda_i)$ , with  $\lambda_i = \exp(-\langle \mathbf{x}_i, \theta_* \rangle)$ , where  $\mathbf{x}_i$  denotes the  $i$ -th row of  $X$ .

4.2.1. *Effect of the initial estimator.* Using a similar experiment as above, we compare the effect of the initial estimator on OLAP. The results is reported on Figure 5. We first note that the Poisson regression requires comparatively larger sample size for the convergence of OLAP. And the behavior of the three estimators are much more similar.

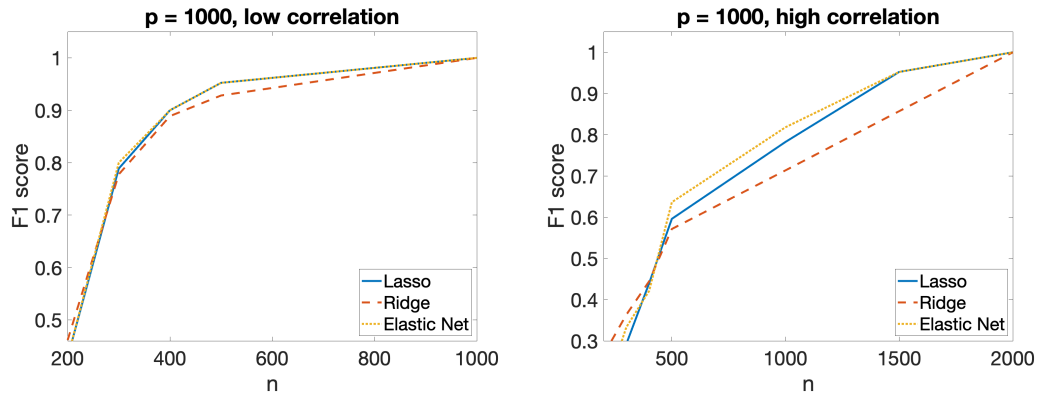


FIGURE 5. Comparison of F1-scores for different initial estimators for Poisson models, low correlation (left) and high-correlation

4.2.2. *Mixing time.* We use the same  $L$ -lag coupling methodology to compare the mixing time of the OLAP Gibbs sampler and the Metropolis-within-Gibbs sampler for the Exact method. Here we set  $n = p = 1000$  in the low correlation case, and  $n = 2000$ ,  $p = 1000$  in the high correlation case.




FIGURE 6. Iterations and mixing time (in seconds) of Poisson regression, low correlation.

Again, We could observe from the Boxplot 7 and 6 that OLAP sampler takes fewer iterations and less running time than Exact sampler to mix, as in the low correlation case, OLAP sampler has a median burn-in iterations before convergence of 54.5, and a median burn-in running time of 0.9 seconds, and EXACT sampler has a median burn-in iterations before convergence of 6186, and a median burn-in running time of 44.3 seconds; while in the high correlation case, OLAP sampler has a median burn-in iterations before convergence of 1319.5, and a median burn-in running time of 25.5 seconds, and EXACT sampler has a median burn-in iterations before convergence of 3480, and a median burn-in running time of 46.8 seconds.




FIGURE 7. Iterations and mixing time (in seconds) of Poisson regression, high correlation.

TABLE 3. F1-score for Poisson regression.  $p = 1000$ , low correlation

n	One-step Lasso		Exact		Lasso	
	Median	Std. Error	Median	Std. Error	Median	Std. Error
200	0.429	0.135	0.533	0.170	0.300	0.080
300	0.789	0.133	0.833	0.131	0.294	0.083
400	0.900	0.107	0.947	0.197	0.280	0.098
500	0.952	0.057	1.000	0.144	0.288	0.097
1000	1.000	0.007	1.000	0.016	0.339	0.122

TABLE 4. F1-score for Poisson regression.  $p = 1000$ , high correlation

n	One-step Lasso		Exact		Lasso	
	Median	Std. Error	Median	Std. Error	Median	Std. Error
200	0.222	0.132	0.286	0.164	0.229	0.064
300	0.293	0.160	0.556	0.189	0.265	0.070
400	0.440	0.168	0.594	0.198	0.281	0.076
500	0.596	0.189	0.778	0.130	0.314	0.057
1000	0.783	0.089	0.894	0.115	0.310	0.052
1500	0.952	0.062	1.000	0.076	0.286	0.046
2000	1.000	0.057	1.000	0.070	0.299	0.066

**4.3. A breast cancer data application.** We illustrate the method with a high-profile breast cancer data example taken from van de Vijver et al. (2002). Accurate prediction of distant metastasis development in breast cancer can help guide treatments, and ultimately save life while preserving quality of life. In van 't Veer (2002), using gene expression data, the authors developed a prognostic score based on 70 identified genes for predicting the advent of metastasis of breast cancer cells in distant organs within 10 years of diagnostics. In this illustration we re-analyzed the data collected by the same research group in van de Vijver et al. (2002) and we compare their 70-genes profile rule with a logistic regression model for predicting the advent of distant metastasis within 5 years.

The data contains 295 women, all younger than 53, with stage I or II breast cancer. Gene intensity measurements of 24496 genes along with 13 clinical variables were collected on the patients. Of the 295 patients, 151 had lymph-node-negative disease, and 144 had lymph-node-positive disease. Ten of the lymph-node-negative, and 120 of the lymph-node-positive had received more aggressive therapy, including chemotherapy, or hormonal therapy, or both. Clearly, the advent of distant metastasis depends on the stage of the cancer when detected, and on initial treatment received. However, despite the cases of lymph-node diseases, most of the cancer cases appear to be at similar stages, and following van de Vijver et al. (2002), we will not account for these interactions in the analysis. As response variable  $y$ , we consider the advent of distant metastasis within five years, as a binary response  $\{0, 1\}$ .

To reduce the size of the covariates, we adopted the pre-processing method implemented by Guo (2018), with one-at-the-time initial logistic regressions that keeps only genes with p-value less than 0.01. To this initial set of genes, we then add the 70 genes identified by van de Vijver et al. (2002), if they are not already selected by the individual T-tests. This pruning process generates a dataset with 295 patients and 1083 genes.

We compare the performance of SparseVB, OLAP, and 70-genes, the predictive model of van de Vijver et al. (2002) based on the 70 genes that they have identified. For the comparison we use a 50-fold cross-validation procedure. In each cross-validation replication the test sample size is 30, and the remaining 265 samples are used for training. To avoid distortion, when sampling the test set, we require the number of 1 and 0 be approximately equal. As a performance metric we compute the  $F_1$ -score in correctly predicting the outcome variable on the test set. Table 5 shows the mean, median and standard deviation of the  $F_1$ -score from the 50 cross-validation replications. The results clearly show a better performance of OLAP both in terms of accuracy and stability.

TABLE 5. F1 score of each algorithm for breast cancer data

	<b>Sparse VB</b>	<b>70-identifiers</b>	<b>OLAP</b>
<b>Mean</b>	0.362	0.593	0.688
<b>Median</b>	0.509	0.609	0.688
<b>Std. Error</b>	0.294	0.109	0.076

**4.4. Mouse PCR data application.** As another illustration, we analyze the mouse PCR data also previously analyzed in Lan et al. (2006); Narisetty et al. (2018). The dataset comprises expression levels of 22575 genes obtained from 31 female and 29 male mice, totaling 60 arrays. Additionally, the physiological phenotype glycerol-3-phosphate acyltransferase (GPAT) was measured using quantitative real-time PCR. These gene expression and phenotypic data are publicly accessible on the GEO database (<http://www.ncbi.nlm.nih.gov/geo>; accession number GSE3330).

We want to predict whether a mouse has low GPAT levels given its genetic expression. The level of GPAT in the body is important, as reduced GPAT levels have been linked to decreased hepatic steatosis, a disease commonly associated with obesity.

Similar to Narisetty et al. (2018), we derive the binary response variable based on GPAT levels, defined as  $y = I(GPAT < Q(0.4))$ , where  $Q(0.4)$  represents the 40% quantile of GPAT. And, given the extensive number of genes, we also did a similar gene pruning as in the previous example, by conducting marginal simple logistic regression of the response  $y$  against individual genes. But unlike Narisetty et al. (2018), and in order to make our experiment more challenging, we select 500 genes that have the most marginally significant p-values, instead of 99 in the original work. Together with the gender variable, this results in  $p = 501$  covariates.

For comparison, we apply SparseVB, S-Gibbs and OLAP. Similar to the breast cancer data experiment, we randomly select 30 different pairs of training and test data sets. To have comparable results to Narisetty et al. (2018), we report in Table 6 the square root of mean squared error (RMSE) as a measure of performance, which is  $RMSE = \sqrt{\frac{1}{n} \sum_i^n (y_i - \hat{\pi}_i)^2}$ , where  $\hat{\pi}_i$  is the probability predicted by the logistic model. We also report in Table 7 the  $F_1$  score in correctly predicting the outcome variable on the test set. The results here again show that OLAP works better than S-Gibbs and SparseVB, in terms of both the accuracy and stability in these measures.

TABLE 6. RMSE for the mouse PCR example

	<b>Sparse VB</b>	<b>Skinny Gibbs</b>	<b>OLAP</b>
<b>Mean</b>	0.4825	0.5751	0.3172
<b>Median</b>	0.5000	0.5236	0.3080
<b>Std. Error</b>	0.0903	0.1279	0.0561

TABLE 7. F1 score for the mouse PCR example

	<b>Sparse VB</b>	<b>Skinny Gibbs</b>	<b>OLAP</b>
<b>Mean</b>	0.4724	0.3786	0.8430
<b>Median</b>	0.5357	0.3636	0.8889
<b>Std. Error</b>	0.3223	0.1993	0.1806

## 5. CONCLUDING REMARKS

Variable selection is an NP-hard problem (Welch (1982)), and algorithms that can solve all versions in polynomial times are unlikely to exist. Therefore, identifying instances of the problem (and corresponding algorithms) that can be solved in polynomial time is of practical importance. Our work in this paper contributes to this literature. We have developed a novel Laplace approximation algorithm that is applicable to a large class of generalized linear models (GLMs) and beyond. The resulting algorithm is fast and accurate, and under mild

conditions, we have shown that it leads to a consistent variable selection methodology in the high-dimensional context. Additionally, we have shown that in many cases the mixing time of the resulting Gibbs sampler scales polynomially in  $(n, p)$ . The simulation results and the real data analysis illustrate the competitiveness of OLAP against some existing high-dimensional variable selection methods. Another advantage of OLAP is that the method has minimal tuning parameter, and as a result can be easily implemented in statistical software.

One important limitation of OLAP is the computational cost of the one-step Newton update. In the current implementation, each component update in the Gibbs sampler is performed at the cost of  $n\|\delta\|_0^2$  operations to form the matrix  $\tilde{\mathcal{H}}^\delta$ , plus  $O(\|\delta\|_0^3/3)$  operations to perform its Cholesky factorization. Finding a recursive update to these calculations will significantly improve the speed of the algorithm. Another potentially useful direction for further investigation is the extension of OLAP beyond GLM, to dealing for instance with hierarchical models, or latent variable models.

#### APPENDIX A. PROOF OF THEOREM 2

*Proof.* We partition the model space  $\Delta$  as

$$\Delta = \bigcup_{\delta_0 \in \mathcal{C}} \Delta(\delta_0),$$

where

$$\mathcal{C} \stackrel{\text{def}}{=} \{\delta \in \Delta : \delta \subseteq \delta_\star\}, \quad \text{and} \quad \Delta(\delta_0) \stackrel{\text{def}}{=} \{\delta' \in \Delta : \delta_0 \subseteq \delta', \text{ and } \min(\delta', \delta_\star) = \delta_0\}.$$

We claim that under the conditions of the theorem the following holds true. For all  $\delta_0 \in \mathcal{C}$ , and all  $M \geq 1$ ,

$$(30) \quad \sum_{\delta \in \Delta(\delta_0)} \frac{\check{\Pi}(\delta|\mathcal{D})}{\check{\Pi}(\delta_0|\mathcal{D})} \leq 2 \quad \text{and} \quad \check{\Pi}(\|\delta\|_0 \geq s_\star + M|\mathcal{D}) \leq 2 \left( \frac{1}{p^{u/2}} \right)^M.$$

To use the claim in (30) to prove the theorem we observe that  $\delta \notin \mathcal{A}_j$  means that either  $\|\delta\|_0 > s_\star + j$ , or  $\min(\delta, \delta_\star)$  is strictly a sub-model of  $\delta_\star$ . We can rewrite this statement as

$$\mathcal{A}_j^c = \{\delta \in \Delta : \|\delta\|_0 > s_\star + j\} \cup \bigcup_{\delta_0 \in \mathcal{C} \setminus \{\delta_\star\}} \{\delta \in \Delta(\delta_0) : \|\delta\|_0 \leq s_\star + j\}.$$

Therefore, using (30),

$$\begin{aligned} \check{\Pi}(\mathcal{A}_j^c|\mathcal{D}) &\leq \frac{2}{p^{u(j+1)/2}} + \check{\Pi}(\delta_\star|\mathcal{D}) \sum_{\delta_0 \in \mathcal{C} \setminus \{\delta_\star\}} \frac{\check{\Pi}(\delta_0|\mathcal{D})}{\check{\Pi}(\delta_\star|\mathcal{D})} \sum_{\delta \in \Delta(\delta_0)} \frac{\check{\Pi}(\delta|\mathcal{D})}{\check{\Pi}(\delta_0|\mathcal{D})} \\ &\leq \frac{2}{p^{u(j+1)/2}} + 2 \sum_{\delta_0 \in \mathcal{C} \setminus \{\delta_\star\}} \frac{\check{\Pi}(\delta_0|\mathcal{D})}{\check{\Pi}(\delta_\star|\mathcal{D})}. \end{aligned}$$

Take  $\delta_0 \in \mathcal{C}$ , with  $\|\delta_0\|_0 = s_\star - k$ , for some  $k > 0$ . Then from (10), we have

$$\frac{\check{\Pi}(\delta_0|\mathcal{D})}{\check{\Pi}(\delta_\star|\mathcal{D})} = p^{uk} \exp(\bar{\ell}^{\delta_0}(\check{\theta}_{\delta_0}; \mathcal{D}) - \bar{\ell}^{\delta_\star}(\check{\theta}_{\delta_\star}; \mathcal{D})).$$

Since the family  $\{\check{\theta}_\delta, \delta \in \Delta\}$  is variable selection consistent,  $\bar{\ell}^{\delta_0}(\check{\theta}^{\delta_0}; \mathcal{D}) - \bar{\ell}^{\delta_\star}(\check{\theta}^{\delta_\star}; \mathcal{D}) \leq -c_2kn$ . It follows that

$$\check{\Pi}(\mathcal{A}_j^c | \mathcal{D}) \leq \frac{2}{p^{u(j+1)/2}} + 2 \sum_{k=1}^{s_\star} \binom{s_\star}{k} \exp(uk \log(p) - c_2kn).$$

Hence for  $c_2n \geq 2(u+1) \log(p)$  as assumed in (14), we obtain

$$\check{\Pi}(\mathcal{A}_j^c | \mathcal{D}) \leq \frac{2}{p^{u(j+1)/2}} + 2 \sum_{k=1}^{s_\star} e^{-c_2nk/2} \leq \frac{2}{p^{u(j+1)/2}} + 4e^{-c_2n/2},$$

which yields the stated inequality.

It remains to prove (30). Given  $\delta_0 \in \mathcal{C}$ ,  $\delta \in \Delta(\delta_0)$ , and setting  $\|\delta\|_0 - \|\delta_0\|_0 = j$ , we have

$$\frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} = \left(\frac{1}{p^u}\right)^j \exp(\bar{\ell}^\delta(\check{\theta}^\delta; \mathcal{D}) - \bar{\ell}^{\delta_0}(\check{\theta}^{\delta_0}; \mathcal{D})).$$

The variable selection consistency of  $\{\check{\theta}_\delta, \delta \in \Delta\}$  yields  $\bar{\ell}^\delta(\check{\theta}^\delta; \mathcal{D}) - \bar{\ell}^{\delta_0}(\check{\theta}^{\delta_0}; \mathcal{D}) \leq c_1j \log(p)$ , so that,

$$(31) \quad \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} \leq \exp(-uj \log(p) + c_1j \log(p)),$$

and we conclude that

$$\begin{aligned} \sum_{\delta \in \Delta(\delta_0)} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} &= \sum_{j \geq 0} \sum_{\delta \in \Delta(\delta_0): \|\delta\|_0 = \|\delta_0\|_0 + j} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} \\ &\leq \sum_{j \geq 0} \binom{p - s_0}{j} e^{-(u-c_1)j \log(p)} \leq \sum_{j \geq 0} e^{-(u-c_1-1)j \log(p)} \leq 2, \end{aligned}$$

provided that  $u \geq 2 + c_1$ , and  $p \geq 2$ , which is satisfied by taking  $u$  as in (13). The second part of (30) follows a similar argument. Since  $\Delta = \cup_{\delta_0 \in \mathcal{C}} \Delta(\delta_0)$ , for any  $M \geq 1$ , we get

$$(32) \quad \check{\Pi}(\|\delta\|_0 \geq s_\star + M | \mathcal{D}) = \sum_{\delta_0 \in \mathcal{C}} \check{\Pi}(\delta_0 | \mathcal{D}) \sum_{\delta \in \Delta(\delta_0): \|\delta\|_0 \geq s_\star + M} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})}.$$

Fix  $\delta_0 \in \mathcal{C}$ , and set  $\|\delta_0\|_0 = s_0$ . For  $k \geq s_\star$ , we have

$$\begin{aligned} \sum_{\delta \in \Delta(\delta_0): \|\delta\|_0 \geq k} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} &= \sum_{j \geq k - s_0} \sum_{\delta \in \Delta(\delta_0): \|\delta\|_0 = s_0 + j} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} \\ &\leq \sum_{j \geq k - s_0} \binom{p - s_0}{j} \exp(-uj \log(p) + c_1j \log(p)) \\ &\leq \sum_{j \geq k - s_0} \exp(-(u - c_1 - 1)j \log(p)). \end{aligned}$$

Hence for  $u/2 \geq (c_1 + 1)$  as assumed in (14) we obtain for  $k > s_0$ , that

$$\sum_{\delta \in \Delta(\delta_0): \|\delta\|_0 \geq k} \frac{\check{\Pi}(\delta | \mathcal{D})}{\check{\Pi}(\delta_0 | \mathcal{D})} \leq \sum_{j \geq k - s_0} \left(\frac{1}{p^{u/2}}\right)^j \leq 2 \left(\frac{1}{p^{u/2}}\right)^{k - s_0}.$$



Hence, the last display together with (32) implies that for any  $M \geq 1$ , we have

$$\check{\Pi}(\|\delta\|_0 \geq s_\star + M|\mathcal{D}) \leq 2 \sum_{\delta_0 \in \mathcal{C}} \check{\Pi}(\delta_0|\mathcal{D}) \left(\frac{1}{p^{u/2}}\right)^M \leq 2 \left(\frac{1}{p^{u/2}}\right)^M.$$

□

## APPENDIX B. PROOF OF THEOREM 6

*Proof.* Let  $q^\delta(\cdot; \mathcal{D})$  denote the quadratic approximation of  $\bar{\ell}^\delta(\cdot; \mathcal{D})$  around  $\tilde{\theta}^\delta$ . Specifically, for  $u \in \mathbb{R}^{\|\delta\|_0}$ ,

$$q^\delta(u; \mathcal{D}) \stackrel{\text{def}}{=} \bar{\ell}^\delta(\tilde{\theta}^\delta; \mathcal{D}) + \left\langle \tilde{\mathcal{G}}_\delta, u - \tilde{\theta}^\delta \right\rangle - \frac{1}{2}(u - \tilde{\theta}^\delta)' \tilde{\mathcal{H}}_\delta (u - \tilde{\theta}^\delta),$$

and

$$\eta^\delta(u; \mathcal{D}) \stackrel{\text{def}}{=} \bar{\ell}^\delta(u; \mathcal{D}) - q^\delta(u; \mathcal{D}).$$

An important step in the argument is to upper bound the remainder  $|\eta^\delta(u; \mathcal{D})|$  for  $u$  close to  $\theta_\star$ . This can be easily done in the setting where the third derivative of  $\psi$  is uniformly bounded. However this will leave out the Poisson model. Following (Bach (2010)), we use the self-concordance assumption in H1-(3) to handle a larger class of link functions  $\psi$ . Specifically, under assumption H1, the following holds. We can find a constant  $C$  that depends only on  $c_3$ ,  $\bar{\kappa}$  and  $b$  in H1 such that for all  $\delta \subseteq \delta_\star$ ,

$$(33) \quad |\eta^\delta(u; \mathcal{D})| \leq C n e^{c_3 b \|u - \tilde{\theta}^\delta\|_1} \times \|u - \tilde{\theta}^\delta\|_1 \times \|u - \tilde{\theta}^\delta\|_2^2, \quad u \in \mathbb{R}^{\|\delta\|_0}.$$

This claim is proved below. Let us assume for the time being that (33) holds, and fix  $\delta, \delta_0 \in \Delta$ . Since  $\hat{\theta}^\delta$  maximizes  $\bar{\ell}^\delta(\cdot; \mathcal{D})$ , we have

$$(34) \quad \bar{\ell}^\delta(\hat{\theta}^\delta; \mathcal{D}) \leq \bar{\ell}^\delta(\tilde{\theta}^\delta; \mathcal{D}).$$

But since  $\tilde{\theta}^\delta$  maximizes  $q^\delta(\cdot; \mathcal{D})$ , we have

$$(35) \quad \begin{aligned} \bar{\ell}^\delta(\tilde{\theta}^\delta; \mathcal{D}) &= q^\delta(\tilde{\theta}^\delta; \mathcal{D}) + \eta^\delta(\tilde{\theta}^\delta; \mathcal{D}) \geq q^\delta(\hat{\theta}^\delta; \mathcal{D}) + \eta^\delta(\tilde{\theta}^\delta; \mathcal{D}) \\ &= \bar{\ell}^\delta(\hat{\theta}^\delta; \mathcal{D}) + \eta^\delta(\tilde{\theta}^\delta; \mathcal{D}) - \eta^\delta(\hat{\theta}^\delta; \mathcal{D}). \end{aligned}$$

We combine (34) and (35) to conclude that

$$(36) \quad \bar{\ell}^\delta(\tilde{\theta}^\delta; \mathcal{D}) - \bar{\ell}^{\delta_0}(\tilde{\theta}^{\delta_0}; \mathcal{D}) \leq \left[ \bar{\ell}^\delta(\hat{\theta}^\delta; \mathcal{D}) - \bar{\ell}^{\delta_0}(\hat{\theta}^{\delta_0}; \mathcal{D}) \right] + |\eta^{\delta_0}(\tilde{\theta}^{\delta_0}; \mathcal{D})| + |\eta^\delta(\tilde{\theta}^\delta; \mathcal{D})|.$$

Hence, proving the theorem boils down to showing that there exists a constant  $C < \infty$  such that for all  $\delta_0 \subseteq \delta_\star$ ,

$$(37) \quad |\eta^{\delta_0}(\tilde{\theta}^{\delta_0}; \mathcal{D})| + |\eta^\delta(\tilde{\theta}^\delta; \mathcal{D})| \leq C \log(p).$$

Indeed, if the MLE family  $\{\hat{\theta}^\delta, \delta \in \Delta\}$  is variable selection consistent with constant  $c_1, c_2$ , say, then for  $\delta_0 \subseteq \delta$ , and  $\min(\delta, \delta_\star) = \delta_0$ , we can conclude from the above argument that

$$\bar{\ell}^\delta(\tilde{\theta}^\delta; \mathcal{D}) - \bar{\ell}^{\delta_0}(\tilde{\theta}^{\delta_0}; \mathcal{D}) \leq c_1 (\|\delta\|_0 - \|\delta_0\|_0) \log(p) + C \log(p) \leq (c_1 + C) (\|\delta\|_0 - \|\delta_0\|_0) \log(p).$$

Similarly, for  $\delta_0 \subseteq \delta_\star$ , applying again (36) with  $\delta_0 \leftarrow \delta_\star$ , and  $\delta \leftarrow \delta_0$ , we get

$$\bar{\ell}^{\delta_\star}(\tilde{\theta}^{\delta_\star}; \mathcal{D}) - \bar{\ell}^{\delta_0}(\tilde{\theta}^{\delta_0}; \mathcal{D}) \geq c_2 (\|\delta_\star\|_0 - \|\delta_0\|_0) n - C \log(p) \geq \frac{c_2}{2} (\|\delta_\star\|_0 - \|\delta_0\|_0) n,$$

for  $c_2 n \geq 2C \log(p)$ . This establishes the theorem. It remains to prove (37) and (33).

Proof of (33). Given,  $u, h \in \mathbb{R}$ , with  $h \neq 0$ , define

$$g(t) = \psi(u + th), \quad t \in [0, 1].$$

H1-(3) implies that  $|g'''(t)| \leq c_3 |h| g''(t)$ , which in turn implies that for all  $t \in [0, 1]$ ,

$$-c_3 |h| \leq \frac{d \log g''(t)}{dt} \leq c_3 |h|.$$

Integrating these two inequalities thrice we obtain for all  $t \in [0, 1]$ ,

$$(38) \quad g''(0) e^{-c_3 |h| t} \leq g''(t) \leq g''(0) e^{c_3 |h| t},$$

$$(39) \quad g''(0) \times \frac{1 - e^{-c_3 |h| t}}{c_3 |h|} \leq g'(t) - g'(0) \leq g''(0) \times \frac{e^{c_3 |h| t} - 1}{c_3 |h|},$$

and

$$-g''(0) \times \frac{1 - c_3 |h| t - e^{-c_3 |h| t}}{c_3^2 h^2} \leq g(t) - g(0) - t g'(0) \leq g''(0) \times \frac{e^{c_3 |h| t} - c_3 |h| t - 1}{c_3^2 h^2}.$$

Subtracting  $g''(0)/t^2$  from all sides we get

$$\begin{aligned} -g''(0) \times \frac{1 - c_3 |h| t + (c_3 |h| t)^2 / 2 - e^{-c_3 |h| t}}{c_3^2 h^2} &\leq g(t) - g(0) - t g'(0) - \frac{t^2}{2} g''(0) \\ &\leq g''(0) \times \frac{e^{c_3 |h| t} - (c_3 |h| t)^2 / 2 - c_3 |h| t - 1}{c_3^2 h^2}. \end{aligned}$$

We apply this with  $t = 1$ , and we use a Taylor of  $e^x$  and  $e^{-x}$  to the third order to conclude that

$$(40) \quad \left| \psi(u + h) - \psi(u) - \psi'(u)h - \frac{h^2}{2} \psi''(u) \right| \leq \frac{c_3}{6} |h|^3 e^{c_3 |h|} \psi''(u).$$

Since for all  $u \in \mathbb{R}^{\delta \|_0}$ ,

$$\begin{aligned} \eta^\delta(u; \mathcal{D}) &= \sum_{i=1}^n \left[ \psi(\langle u, \mathbf{x}_i \rangle) - \psi(\langle \tilde{\theta}^\delta, \mathbf{x}_i \rangle) - \psi'(\langle \tilde{\theta}^\delta, \mathbf{x}_i \rangle) \langle u - \tilde{\theta}^\delta, \mathbf{x}_i \rangle \right. \\ &\quad \left. - \frac{1}{2} \psi''(\langle \tilde{\theta}^\delta, \mathbf{x}_i \rangle) \langle u - \tilde{\theta}^\delta, \mathbf{x}_i \rangle^2 \right]. \end{aligned}$$

We conclude with (40) that

$$\begin{aligned} |\eta^\delta(u; \mathcal{D})| &\leq \frac{c_3 b}{6} \|u - \tilde{\theta}^\delta\|_1 e^{c_3 b \|u - \tilde{\theta}^\delta\|_1} \sum_{i=1}^n \psi''(\langle \tilde{\theta}^\delta, \mathbf{x}_i \rangle) \langle u - \tilde{\theta}^\delta, \mathbf{x}_i \rangle^2 \\ &\leq \frac{c_3 b \bar{\kappa} n}{6} e^{c_3 b \|u - \tilde{\theta}^\delta\|_1} \|u - \tilde{\theta}^\delta\|_1 \times \|u - \tilde{\theta}^\delta\|_2^2, \end{aligned}$$

which is the claim (33).

Proof of (37). Fix  $\delta_0 \subseteq \delta_*$ . Since  $\hat{\theta}^{\delta_0} - \tilde{\theta}^{\delta_0} = \hat{\theta}^{\delta_0} - \theta_*^{\delta_0} + \theta_*^{\delta_0} - \tilde{\theta}^{\delta_0}$ , and  $\delta_0 \subseteq \delta_*$ , we get, using (23),

$$\|\hat{\theta}^{\delta_0} - \tilde{\theta}^{\delta_0}\|_2 \leq \|\hat{\theta}^{\delta_*} - \theta_*\|_2 + \|\tilde{\theta} - \theta_*\|_2 \leq C\sqrt{\frac{s_* \log(p)}{n}}.$$

Hence

$$\|\hat{\theta}^{\delta_0} - \tilde{\theta}^{\delta_0}\|_1 \leq C\sqrt{\frac{s_*^2 \log(p)}{n}} \leq C',$$

for  $n \geq s_*^2 \log(p)$ . As a result, using (33), we conclude that

$$|\eta^{\delta_0}(\hat{\theta}^{\delta_0}; \mathcal{D})| \leq Cn\sqrt{\frac{s_*^2 \log(p)}{n}} \times \frac{s_* \log(p)}{n} \leq C' \log(p),$$

under the sample size condition (24). By the definition of  $\check{\theta}^\delta$  in (7),

$$\check{\theta}^{\delta_0} - \tilde{\theta}^{\delta_0} = (\tilde{\mathcal{H}}^{\delta_0})^{-1} \tilde{\mathcal{G}}^{\delta_0},$$

and under H1-(4), the smallest eigenvalue of  $\tilde{\mathcal{H}}^{\delta_0}$  is at least  $\underline{\kappa}n$ . Therefore, and using the expression of the gradient  $\mathcal{G}^\delta$ , we have

$$\|\check{\theta}^{\delta_0} - \tilde{\theta}^{\delta_0}\|_2 \leq \frac{1}{\underline{\kappa}n} \left( \|\mathcal{G}_* \|_{\delta_0} \|_2 + \|\tilde{\mathcal{G}}^{\delta_0} - \mathcal{G}_* \|_{\delta_0} \|_2 \right),$$

where we define  $\mathcal{G}_* \stackrel{\text{def}}{=} \nabla \bar{\ell}^\delta(u; \mathcal{D})|_{u=\theta_*}$ . According to H1-(2),  $\|\mathcal{G}_*\|_\infty \leq M_0 \sqrt{n \log(p)}$ . Hence

$$\frac{\|\mathcal{G}_* \|_{\delta_0} \|_2}{\underline{\kappa}n} \leq C\sqrt{\frac{\|\delta_0\|_0 \log(p)}{n}} \leq C\sqrt{\frac{s_* \log(p)}{n}}.$$

On the other hand, a comparison between the  $j$ -th component of  $\tilde{\mathcal{G}}^{\delta_0}$  and  $[\mathcal{G}_*]_{\delta_0}$  gives

$$\begin{aligned} (\tilde{\mathcal{G}}^{\delta_0})_j - [\mathcal{G}_*]_j &= - \sum_{i=1}^n \left( \psi'(\langle \tilde{\theta}^{\delta_0}, \mathbf{x}_i \rangle) - \psi'(\langle \theta_*, \mathbf{x}_i \rangle) \right) [\mathbf{x}_i]_j \\ &= - \sum_{k: \delta_{*,k}=1} (\tilde{\theta}_k - \theta_{*,k}) \sum_{i=1}^n \psi''(\langle \vartheta, \mathbf{x}_i \rangle) \mathbf{x}_{ij} \mathbf{x}_{ik}, \end{aligned}$$

for some  $\vartheta$  on the segment between  $\tilde{\theta}^{\delta_0}$  and  $\theta_*$ . We can then appeal to H1-(4-5) to conclude that

$$\left| (\tilde{\mathcal{G}}^{\delta_0})_j - [\mathcal{G}_*]_j \right| \leq \bar{\kappa}n |\tilde{\theta}_j - \theta_{*,j}| + M \|\tilde{\theta} - \theta_*\|_1 \sqrt{n \log(p)} \leq \bar{\kappa}n |\tilde{\theta}_j - \theta_{*,j}| + C s_* \log(p).$$

We deduce that

$$\frac{1}{\underline{\kappa}n} \|\tilde{\mathcal{G}}^{\delta_0} - [\mathcal{G}_*]_{\delta_0}\|_2 \leq \frac{Cn}{\underline{\kappa}n} \sqrt{\frac{s_* \log(p)}{n}} + \frac{C s_*^{3/2} \log(p)}{n} \leq C' \sqrt{\frac{s_* \log(p)}{n}}.$$

Therefore, (33), and a similar argument yields

$$|\eta^{\delta_0}(\hat{\theta}^{\delta_0}; \mathcal{D})| \leq Cn\sqrt{\frac{s_*^2 \log(p)}{n}} \times \frac{s_* \log(p)}{n} \leq C' \log(p),$$

which concludes the proof of (37).  $\square$

## APPENDIX C. PROOF OF THEOREM 4

When  $J = 1$ , Algorithm 1 is also known as a random scan Gibbs sampler. It generates a reversible and positive Markov chain with invariant distribution  $\check{\Pi}$  that fits in the framework presented in Section 3. We denote  $K$  its transition kernel. We prove the result by applying Lemma 7 and Proposition 9. The initial distribution of Algorithm 1 has a density  $f_0$  with respect to  $\check{\Pi}$  given by:  $f_0(\delta) = 1/\check{\Pi}(\delta^{(0)}|\mathcal{D})$ , if  $\delta = \delta^{(0)}$ , and zero everywhere else. Hence

$$\|f_0\|_\infty = \frac{1}{\check{\Pi}(\delta^{(0)}|\mathcal{D})} \leq p^\alpha.$$

Fix  $\Delta_0 \subseteq \Delta$  to be determined later such that  $\check{\Pi}(\Delta_0|\mathcal{D}) \geq 15/16$ , and set  $\zeta = 8(1 - \check{\Pi}(\Delta_0|\mathcal{D}))$ . If we can build canonical paths on  $\Delta_0$  such that H2 holds, then by Proposition 9, we get  $\lambda_\zeta(K) \geq 1/m(\Delta_0)$ . And since  $\log(1 - x) \leq -x$  for all  $x \in (0, 1]$ , it follows that for  $N \geq 1$  large enough such that

$$N \geq 2m(\Delta_0) \log\left(\frac{\|f_0\|_\infty^2}{\zeta_0^2}\right),$$

we have

$$\left(1 - \frac{\lambda_\zeta(K)}{2}\right)^N = \exp\left(N \log\left(1 - \frac{\lambda_\zeta(K)}{2}\right)\right) \leq \exp\left(-\frac{N}{2m(\Delta_0)}\right) \leq \frac{\zeta_0^2}{\|f_0\|_\infty^2}.$$

Therefore by (26) of Lemma 7,

$$\|K^N(\delta^{(0)}, \cdot) - \check{\Pi}\|_{\text{tv}}^2 \leq \max\left(\zeta\|f_0\|_\infty^2, \left(1 - \frac{\lambda_\zeta(K)}{2}\right)^N \|f_0\|_\infty^2\right) \leq \max(\zeta\|f_0\|_\infty^2, \zeta_0^2).$$

We apply this result with the choice  $\Delta_0 = \{\delta \in \Delta : \|\delta\|_0 \leq s_\star + J_0\}$ . By (30),

$$\zeta\|f_0\|_\infty^2 = 8\|f_0\|_\infty^2(1 - \check{\Pi}(\Delta_0|\mathcal{D})) \leq \frac{16\|f_0\|_\infty^2}{p^{u(J_0+1)/2}} \leq \frac{16}{p^{u/2}} \frac{\|f_0\|_\infty^2}{p^{uJ_0/2}} \leq \frac{16}{p^{u/2}},$$

by taking  $J_0 = 4\alpha/u$ . Hence, for

$$N \geq 4m(\Delta_0) (\log(1/\zeta_0) + \alpha \log(p)),$$

$$(41) \quad \|K^N(\delta^{(0)}, \cdot) - \check{\Pi}\|_{\text{tv}} \leq \max\left(\frac{4}{p^{u/4}}, \zeta_0\right).$$

We now use the canonical path argument to upper bound the term  $m(\Delta_0)$ . First, we build a graph on  $\Delta_0$  by putting an edge between  $\delta$  and  $\delta'$  if  $\|\delta - \delta'\|_0 = 1$ . Clear H2 holds. We build the canonical paths by removing or adding variables as follow. For any  $\delta$  we first build a path between  $\delta$  and  $\delta_\star$ . First, we set to zero (1 term at the time and in their decreasing index order) the components  $j$  of  $\delta$  for which  $\delta_j = 1$  and  $\delta_{\star j} = 0$ . We then reach  $\underline{\delta} \stackrel{\text{def}}{=} \min(\delta, \delta_\star)$ . This corresponds to removing one-by-one non-relevant variables from the model  $\delta$ . Then we set to one (one at the time, and in their increasing index order) the component  $j$  of  $\underline{\delta}$  for which  $\underline{\delta}_j = 0$ , and  $\delta_{\star j} = 1$ . We then reach  $\delta_\star$ . This corresponds to adding one-by-one relevant variables not already contained in  $\underline{\delta}$ .

Given two arbitrary points  $\delta$  and  $\delta'$ , we build the canonical path between them as follows. Let  $\vartheta$  denote the point where the canonical path from  $\delta$  to  $\delta_\star$  and the canonical path from

$\delta'$  to  $\delta_\star$  meet for the first time. The canonical path from  $\delta$  to  $\delta'$  is obtained by following the canonical path from  $\delta$  to  $\delta_\star$  until  $\vartheta$  is reached, then we follow the canonical path from  $\delta'$  to  $\delta_\star$  in reverse direction starting from  $\vartheta$  until  $\delta'$ . We have

$$(42) \quad |\gamma_{\delta, \delta'}| \leq 2(s_\star + J_0).$$

Let  $x, y$  denote generic elements of  $\Delta_0$ . Fix an edge  $e = (\delta', \delta)$ . Let  $\Lambda(e)$  be the set of all elements of  $\Delta_0$  whose path to  $\delta_\star$  go through  $e$ . If  $\gamma_{xy} \ni e$  then  $x \in \Lambda(e)$  or  $y \in \Lambda(e)$  (but not both). Therefore,

$$\begin{aligned} \sum_{\gamma_{xy}: \gamma_{xy} \ni e} \frac{|\gamma_{xy}| \check{\Pi}(x|\mathcal{D}) \check{\Pi}(y|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D}) K(\delta', \delta)} &\leq 2 \sum_{x \in \Lambda(e)} \sum_{y \in \mathcal{X}_0} \frac{|\gamma_{xy}| \check{\Pi}(x|\mathcal{D}) \check{\Pi}(y|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D}) K(\delta', \delta)} \\ &\leq 2 \sum_{x \in \Lambda(e)} \frac{|\gamma_{xy}| \check{\Pi}(x|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D}) K(\delta', \delta)}. \end{aligned}$$

Using the last display with (42), we conclude that

$$\mathfrak{m}(\Delta_0) \leq 4(s_\star + J_0) \max_{(\delta', \delta) \in \mathcal{E}_0} \frac{1}{K(\delta', \delta)} \sum_{x \in \Lambda(e)} \frac{\Pi(x|\mathcal{D})}{\Pi(\delta'|\mathcal{D})}.$$

**Case 1:**  $\delta = \delta^{(j,0)}$ , and  $\delta' = \delta^{(j,1)}$ , for some  $j$  such that  $\delta_{\star j} = 0$ . In this case,

$$\begin{aligned} K(\delta', \delta) &= \frac{1}{p} (1 - q_j(\delta')) \\ &= \frac{1}{p} \left( 1 + \exp(-\mathbf{u} \log(p) + \bar{\ell}^{\delta^{(j,1)}}(\check{\theta}_{\delta^{(j,1)}}; \mathcal{D}) - \bar{\ell}^{\delta^{(j,0)}}(\check{\theta}_{\delta^{(j,0)}}; \mathcal{D})) \right)^{-1} \\ &\geq \frac{1}{p} \left( 1 + \frac{1}{p^{u/2}} \right)^{-1} \geq \frac{1}{p} \left( 1 - \frac{1}{p^{u/2}} \right) \geq \frac{1}{2p}, \end{aligned}$$

where we use (11) to claim that  $\bar{\ell}^{\delta^{(j,1)}}(\check{\theta}_{\delta^{(j,1)}}; \mathcal{D}) - \bar{\ell}^{\delta^{(j,0)}}(\check{\theta}_{\delta^{(j,0)}}; \mathcal{D}) \leq c_1 \log(p)$ , and for  $\mathbf{u} > 2c_1$ , and  $p^{u/2} \geq 2$ . Furthermore, in this case, note that  $\delta' \supset \delta$ , and differs from  $\delta$  only at some  $j$  such that  $\delta_{\star j} = 0$ . Therefore,  $\Lambda(e) = \Delta(\delta')$ , that is the elements  $\vartheta \in \Delta_{s_\star + J_0}$  contains  $\delta'$  and differs from  $\delta'$  only at components at which  $\delta_{\star k} = 0$ . Therefore, as seen in the proof of Theorem 2,

$$\sum_{x \in \Lambda(e)} \frac{\check{\Pi}(x|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D})} \leq 2.$$

and it follows that

$$\frac{1}{K(\delta', \delta)} \sum_{x \in \Lambda(e)} \frac{\Pi(x|\mathcal{D})}{\Pi(\delta'|\mathcal{D})} \leq 4p.$$

**Case 2:**  $\delta' \subset \delta_*$ ,  $\delta' = \delta^{(j,0)}$ , and  $\delta = \delta^{(j,1)}$ , for some  $j$  such that  $\delta_{*j} = 1$ . In this case, using (12),

$$\begin{aligned} K(\delta', \delta) &= \frac{q_j(\delta')}{p} = \frac{1}{p} \left( 1 + \exp(\mathfrak{u} \log(p) - \bar{\ell}^{\delta^{(j,1)}}(\check{\theta}_{\delta^{(j,1)}}; \mathcal{D}) + \bar{\ell}^{\delta^{(j,0)}}(\check{\theta}_{\delta^{(j,0)}}; \mathcal{D})) \right)^{-1} \\ &\geq \frac{1}{p} (1 + \exp(\mathfrak{u} \log(p) - c_2 n))^{-1} \geq \frac{1}{p} \left( 1 + \frac{1}{e^{c_2 n/2}} \right) \geq \frac{1}{2p}, \end{aligned}$$

for  $c_2 n \geq 2\mathfrak{u} \log(p)$ . In this configuration, we see that

$$\Lambda(e) = \bigcup_{\vartheta \subseteq \delta'} \Delta(\vartheta),$$

therefore,

$$\sum_{x \in \Lambda(e)} \frac{\check{\Pi}(x|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D})} = \sum_{\vartheta \subseteq \delta'} \frac{\check{\Pi}(\vartheta|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D})} \sum_{x \in \Delta(\vartheta)} \frac{\check{\Pi}(x|\mathcal{D})}{\check{\Pi}(\vartheta|\mathcal{D})} \leq 2 \sum_{\vartheta \subseteq \delta'} \frac{\check{\Pi}(\vartheta|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D})}.$$

We proceed similarly as in the proof of Theorem 2 to show that

$$\sum_{\vartheta \subseteq \delta'} \frac{\check{\Pi}(\vartheta|\mathcal{D})}{\check{\Pi}(\delta'|\mathcal{D})} \leq 1 + 4e^{-c_2 n/2} \leq 2,$$

for  $c_2 n \geq 2(\mathfrak{u} + 1) \log(p)$ . We conclude that for some absolute constant  $C$ .

$$\mathfrak{m}(\Delta_0) \leq C(s_* + J_0)p.$$

We combine this bound with (41) and  $J_0 = 4\alpha/\mathfrak{u}$ , to conclude that

$$\|K^N(\delta^{(0)}, \cdot) - \check{\Pi}\|_{\text{tv}} \leq \max\left(\frac{4}{p^{\mathfrak{u}/4}, \zeta_0\right), \text{ for } N \geq C\left(s_* + \frac{4\alpha}{\mathfrak{u}}\right) (\log(1/\zeta_0) + \alpha \log(p))p,$$

for some absolute constant  $C$ . This concludes the proof.

#### APPENDIX D. PROOF OF THE RESULTS OF SECTION 3

**D.1. Proof of Lemma 7.** By reversibility, for all  $j \geq 1$ , the density of  $\pi_0 K^j$  with respect to  $\pi$  is  $f_j = K^j f_0$ . Therefore

$$\|\pi_0 K^j - \pi\|_{\text{tv}} = \int |f_j(x) - 1| \pi(dx) \leq \sqrt{\text{Var}_\pi(f_j)}.$$

Take  $f \in L^2(\pi)$ . Since  $\pi(f) = \pi(Kf)$ , and  $\text{Var}_\pi(f) = \langle f, f \rangle_\pi - \pi(f)^2$ , we have

$$\text{Var}_\pi(Kf) - \text{Var}_\pi(f) = \langle Kf, Kf \rangle_\pi - \langle f, f \rangle_\pi = \langle f, K_* Kf \rangle_\pi - \langle f, f \rangle_\pi = -\mathcal{E}_{K_* K}(f, f).$$

If  $K$  is positive (it is therefore  $\pi$ -reversible), then  $K_* K = K^2$ , and  $K$  admits a square root: there exists a bounded  $\pi$ -reversible operator  $S$  such that  $S^2 = K$ , and  $S$  commutes with  $K$ . Furthermore, with  $\mathbb{I}$  denoting the identity operator,  $\mathbb{I} - K$  is also a positive operator:  $\mathbb{I} - K$  is clearly  $\pi$ -reversible, and  $\langle f, (\mathbb{I} - K)f \rangle_\pi = \|f\|_2^2 - \langle f, Kf \rangle_\pi \geq 0$ , using the fact that the operator norm of  $K$  is smaller or equal to one. Hence

$$\langle f, Kf \rangle_\pi - \langle f, K_* Kf \rangle_\pi = \langle f, (K - K^2)f \rangle_\pi = \langle f, S(\mathbb{I} - K)Sf \rangle_\pi = \langle Sf, (\mathbb{I} - K)Sf \rangle_\pi \geq 0.$$

Therefore when  $K$  is positive we can have

$$\mathrm{Var}_\pi(Kf) \leq \mathrm{Var}_\pi(f) - \mathcal{E}_K(f, f).$$

In particular, for all  $j \geq 1$ ,

$$(43) \quad \mathrm{Var}_\pi(f_j) \leq \mathrm{Var}_\pi(f_{j-1}) - \mathcal{E}_K(f_{j-1}, f_{j-1}).$$

First, we observe that the last display implies that  $\{\mathrm{Var}_\pi(f_k), k \geq 0\}$  is non-increasing. Fix  $j \geq 1$ , and suppose now that  $\mathrm{Var}_\pi(f_j) > \zeta \|f_0\|_\infty^2$ . Since the sequence  $\{\mathrm{Var}_\pi(f_k), k \geq 0\}$  is non-increasing, for all  $0 \leq i \leq j$ ,  $\mathrm{Var}_\pi(f_i) > \zeta \|f_0\|_\infty^2$ . Note also that  $\|f_k\|_\infty \leq \|f_0\|_\infty$  for all  $k \geq 0$ . Therefore, for all  $1 \leq i \leq j$ ,

$$\begin{aligned} \mathrm{Var}_\pi(f_i) &\leq \mathrm{Var}_\pi(f_{i-1}) - \mathcal{E}_K \left( \frac{f_{i-1}}{\|f_0\|_\infty}, \frac{f_{i-1}}{\|f_0\|_\infty} \right) \|f_0\|_\infty^2 \\ &\leq \mathrm{Var}_\pi(f_{i-1}) - \|f_0\|_\infty^2 \lambda_\zeta(K) \left( \mathrm{Var}_\pi \left( \frac{f_{i-1}}{\|f_0\|_\infty} \right) - \frac{\zeta}{2} \right), \\ &\leq \mathrm{Var}_\pi(f_{i-1}) - \frac{\lambda_\zeta(K)}{2} \mathrm{Var}_\pi(f_{i-1}) \\ &\leq \left( 1 - \frac{\lambda_\zeta(K)}{2} \right)^i \mathrm{Var}_\pi(f_0). \end{aligned}$$

We conclude that for all  $j \geq 1$ ,

$$\mathrm{Var}_\pi(f_j) \leq \max \left[ \zeta \|f_0\|_\infty^2, \left( 1 - \frac{\lambda_\zeta(K)}{2} \right)^j \mathrm{Var}_\pi(f_0) \right].$$

This ends the proof.

**D.2. Proof of Theorem 8.** The proof follows the same argument in the proof of Cheeger's inequality due to Lawler and Sokal (1988). Fix a measurable set  $A$  such that  $\epsilon < \pi(A) < 1 - \epsilon$ , and set

$$f_A(\cdot) = \pi(A)\mathbf{1}_{A^c}(\cdot) - \pi(A^c)\mathbf{1}_A(\cdot).$$

Clearly  $\pi(f_A) = 0$ , and we check that  $(I - K)f_A(\cdot) = \mathbf{1}_{A^c}(\cdot) - K(\cdot, A^c)$ , and

$$\mathcal{E}_K(f_A, f_A) = \langle f_A, (I - K)f_A \rangle_\pi = \int_A \pi(dx) K(x, A^c).$$

Also we have

$$\mathrm{Var}_\pi(f_A) = \pi(A)\pi(A^c) = \pi(A)(1 - \pi(A)).$$

We check that for all  $\epsilon \in [0, 1/2)$ ,

$$\frac{1 - \sqrt{1 - 2\epsilon}}{2} < \epsilon < 1 - \epsilon < \frac{1 + \sqrt{1 - 2\epsilon}}{2},$$

which implies that for all  $\epsilon \leq x \leq 1 - \epsilon$ ,  $x(1 - x) > \epsilon/2$ . We use this to conclude that  $\mathrm{Var}_\pi(f_A) = \pi(A)(1 - \pi(A)) > \epsilon/2$ . Furthermore,

$$(\pi(A) - \epsilon)(\pi(A^c) - \epsilon) = \pi(A)\pi(A^c) - \epsilon + \epsilon^2 \leq \pi(A)\pi(A^c) - \frac{\epsilon}{2} \leq \mathrm{Var}_\pi(f_A) - \frac{\epsilon}{4}.$$

It follows that

$$\frac{\int_A \pi(dx) K(x, A^c)}{(\pi(A) - \epsilon)(\pi(A^c) - \epsilon)} \geq \frac{\mathcal{E}_K(f_A, f_A)}{\mathbf{Var}_\pi(f_A) - \frac{\epsilon}{4}} \geq \lambda_{\epsilon/2}(K).$$

We use this to conclude as claimed that

$$\lambda_{\epsilon/2}(K) \leq \Phi_\epsilon(K).$$

Take  $f \in L^2(\pi)$  with  $\|f\|_\infty \leq 1$ , and  $\mathbf{Var}_\pi(f) > \epsilon$ . We aim to lower bound the term  $\mathcal{E}_K(f, f)/(\mathbf{Var}_\pi(f) - \epsilon/2)$ . Given  $c \in \mathbb{R}$ , let  $f_c \stackrel{\text{def}}{=} f - c$ , and let  $\bar{K}(dx, dy) \stackrel{\text{def}}{=} \pi(dx)K(x, dy)$ . By the Cauchy-Schwarz inequality

$$\begin{aligned} & \left( \int \bar{K}(dx, dy) |f_c^2(y) - f_c^2(x)| \right)^2 \\ & \leq \int \bar{K}(dx, dy) (f_c(y) - f_c(x))^2 \times \int \bar{K}(dx, dy) (f_c(y) + f_c(x))^2 \\ & \leq 4\pi(f_c^2) \int \bar{K}(dx, dy) (f(y) - f(x))^2. \end{aligned}$$

We deduce that

$$\sqrt{\mathcal{E}_K(f, f)} \geq \frac{\int \bar{K}(dx, dy) |f_c^2(y) - f_c^2(x)|}{\sqrt{8\pi(f_c^2)}}.$$

By reversibility,

$$\int \bar{K}(dx, dy) |f_c^2(y) - f_c^2(x)| = 2 \int_{\{(x, y): f_c^2(y) > f_c^2(x)\}} \bar{K}(dx, dy) (f_c^2(y) - f_c^2(x)).$$

For  $\alpha \geq 0$ , we set  $A_\alpha \stackrel{\text{def}}{=} \{x : f_c^2(x) > \alpha\}$ , and  $\bar{A}_\alpha \stackrel{\text{def}}{=} \{x : f_c^2(x) \leq \alpha\}$  the complement of  $A_\alpha$ . Let

$$I_c \stackrel{\text{def}}{=} [\mathbf{ess-inf}(f_c^2), \mathbf{ess-sup}(f_c^2)],$$

We note that  $\pi(A_\alpha) = 0$  for  $\alpha > \mathbf{ess-sup}(f_c^2)$ , and for  $\alpha < \mathbf{ess-inf}(f_c^2)$ ,  $\pi(\bar{A}_\alpha) = 0$ . Hence in both cases  $\int_{A_\alpha} \pi(dx)K(x, \bar{A}_\alpha) = 0$ . We also observe that for all  $x, y \in \mathcal{X}$ ,

$$\int \mathbf{1}_{A_\alpha}(x) \mathbf{1}_{\bar{A}_\alpha}(y) d\alpha = (f_c^2(x) - f_c^2(y)) \mathbf{1}_{\{f_c^2(x) > f_c^2(y)\}}.$$

Using this and Fubini's theorem,

$$\begin{aligned} \int \bar{K}(dx, dy) |f_c^2(y) - f_c^2(x)| &= 2 \int d\alpha \int \bar{K}(dx, dy) \mathbf{1}_{A_\alpha}(x) \mathbf{1}_{\bar{A}_\alpha}(y) \\ &= 2 \int_{I_c} d\alpha \int_{A_\alpha} \pi(dx) K(x, \bar{A}_\alpha) \\ &\geq 2\Phi_{\epsilon'}(K) \int_{I_c} (\pi(A_\alpha) - \epsilon') (\pi(\bar{A}_\alpha) - \epsilon') d\alpha \\ &= 2\Phi_{\epsilon'}(K) \int_{I_c} (\pi(A_\alpha)\pi(\bar{A}_\alpha) - \epsilon'(1 - \epsilon')) d\alpha, \end{aligned}$$



where  $\epsilon' \stackrel{\text{def}}{=} \epsilon/32$ .

$$\begin{aligned} \int_{I_c} \pi(A_\alpha) \pi(\bar{A}_\alpha) d\alpha &= \int_0^\infty \pi(A_\alpha) \pi(\bar{A}_\alpha) d\alpha \\ &= \int (f_c^2(y) - f_c^2(x)) \mathbf{1}_{\{(x,y): f_c^2(y) > f_c^2(x)\}} \pi(dx) \pi(dy) \\ &= \frac{1}{2} \int |f_c^2(y) - f_c^2(x)| \pi(dx) \pi(dy), \end{aligned}$$

whereas

$$\int_{I_c} \epsilon'(1 - \epsilon') d\alpha = \epsilon'(1 - \epsilon') (\text{ess-sup}(f_c^2) - \text{ess-inf}(f_c^2)).$$

For  $y, z \in \mathcal{X}$ ,

$$|f_c(y)|^2 - |f_c(z)|^2 \leq |f(y) - f(z)| |f_c(y) + f_c(z)| \leq 4\|f\|_\infty \|f_c\|_\infty,$$

which implies that

$$\int_{I_c} \epsilon'(1 - \epsilon') d\alpha \leq 4\epsilon' \|f_c\|_\infty.$$

We conclude that

$$(44) \quad \sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{8\pi(f_c^2)}} \left( \int |f_c^2(y) - f_c^2(x)| \pi(dx) \pi(dy) - 8\epsilon' \|f_c\|_\infty \right).$$

Since (44) holds true for all  $c$ , we get, by dominated convergence that

$$\begin{aligned} \sqrt{\mathcal{E}_K(f, f)} &\geq \limsup_{c \rightarrow +\infty} \frac{\Phi_{\epsilon'}(K)}{\sqrt{8\pi(f_c^2)}} \left( \int |f_c^2(y) - f_c^2(x)| \pi(dx) \pi(dy) - 8\epsilon' \|f_c\|_\infty \right) \\ &= \frac{\Phi_{\epsilon'}(K)}{\sqrt{8}} \left( 2 \int |f(y) - f(x)| \pi(dx) \pi(dy) - 8\epsilon' \right). \end{aligned}$$

By Jensen's inequality

$$\int \int |f(y) - f(x)| \pi(dx) \pi(dy) \geq \int |f(x) - \pi(f)| \pi(dx) = \pi(|\bar{f}|),$$

where  $\bar{f} \stackrel{\text{def}}{=} f - \pi(f)$ . If  $\pi(|\bar{f}|) \geq \sqrt{\text{Var}_\pi(f)}/2$ , we conclude that

$$\sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{8}} \left( \sqrt{\text{Var}_\pi(f)} - 8\epsilon' \right).$$

Since the function  $f$  is taken such that  $\text{Var}_\pi(f) > \epsilon$ , using the convexity inequality  $\sqrt{x+y} \geq (\sqrt{x} + \sqrt{y})/\sqrt{2}$  valid for all  $x, y \geq 0$  we have

$$\begin{aligned} \sqrt{\text{Var}_\pi(f)} - 8\epsilon' &= \sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2} + \frac{\epsilon}{2}} - \frac{\epsilon}{4} \geq \frac{\sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2}}}{\sqrt{2}} + \frac{\epsilon^{1/2}}{2} - \frac{\epsilon}{4} \\ &\geq \frac{\sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2}}}{\sqrt{2}}. \end{aligned}$$

It follows that if  $\pi(|\bar{f}|) \geq \sqrt{\text{Var}_\pi(f)}/2$ , we have

$$(45) \quad \sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{16}} \sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2}}.$$

Going back to (44), using the value at  $c = \pi(f)$ , and recalling the notation  $\bar{f} = f - \pi(f)$ , we have

$$(46) \quad \sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{8\pi(\bar{f}^2)}} \left( \int \int |\bar{f}^2(y) - \bar{f}^2(x)| \pi(dx)\pi(dy) - 16\epsilon' \right).$$

We have

$$\int |\bar{f}^2(y) - \bar{f}^2(x)| \pi(dy) \geq \left| \int (\bar{f}^2(y) - \bar{f}^2(x)) \pi(dy) \right| = |\bar{f}^2(x) - \text{Var}_\pi(f)|,$$

so that

$$\begin{aligned} \int \int |\bar{f}^2(y) - \bar{f}^2(x)| \pi(dx)\pi(dy) &\geq \int |\bar{f}^2(x) - \text{Var}_\pi(f)| \pi(dx) \\ &= \int \{ \bar{f}^2(x) + \text{Var}_\pi(f) - 2 \min(\bar{f}^2(x), \text{Var}_\pi(f)) \} \pi(dx) \\ &= 2\text{Var}_\pi(f) - 2\text{Var}_\pi(f) \int \min\left( \left( \frac{\bar{f}(x)}{\sqrt{\text{Var}_\pi(f)}} \right)^2, 1 \right) \pi(dx) \\ &\geq 2\text{Var}_\pi(f) - 2\text{Var}_\pi(f) \frac{\pi(|\bar{f}|)}{\sqrt{\text{Var}_\pi(f)}}, \end{aligned}$$

where the last inequality in the last display follows from the observation that  $\min(1, x^2) \leq |x|$  for all  $x \in \mathbb{R}$ . Hence if  $\pi(|\bar{f}|) \leq \sqrt{\text{Var}_\pi(f)}/2$  we have

$$\int \int |f^2(y) - f^2(x)| \pi(dx)\pi(dy) \geq \text{Var}_\pi(f),$$

and we use this with (46) to deduce that

$$\sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{8}} \times \left( \sqrt{\text{Var}_\pi(f)} - \frac{16\epsilon'}{\sqrt{\text{Var}_\pi(f)}} \right).$$

Since  $\text{Var}_\pi(f) > \epsilon$ ,

$$\begin{aligned} \sqrt{\text{Var}_\pi(f)} - \frac{16\epsilon'}{\sqrt{\text{Var}_\pi(f)}} &\geq \sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2} + \frac{\epsilon}{2} - \frac{\epsilon^{1/2}}{2}} \\ &\geq \frac{\sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2}}}{\sqrt{2}} + \frac{\epsilon^{1/2}}{2} - \frac{\epsilon^{1/2}}{2} \geq \frac{\sqrt{\text{Var}_\pi(f) - \frac{\epsilon}{2}}}{\sqrt{2}}. \end{aligned}$$

In conclusion, if  $\text{Var}_\pi(f) > \epsilon$  we have

$$(47) \quad \sqrt{\mathcal{E}_K(f, f)} \geq \frac{\Phi_{\epsilon'}(K)}{\sqrt{16}} \left( \sqrt{\pi(f^2) - \epsilon/2} \right).$$

We combine the last display and (45) to conclude that for all  $f \in L^2(\pi)$  such that  $\|f\|_\infty \leq 1$ , and  $\text{Var}_\pi(f) > \epsilon$ , we have

$$\frac{\mathcal{E}_K(f, f)}{\pi(f^2) - \epsilon/2} \geq \frac{\Phi_{\epsilon'}^2(K)}{16},$$

which implies the lower bound.

**D.3. Proof of Proposition 9.** Take a function  $f : \mathcal{X} \rightarrow \mathbb{R}$  such that  $\|f\|_\infty \leq 1$ , and  $\text{Var}_\pi(f) > \zeta$ . We have

$$\begin{aligned} 2\text{Var}_\pi(f) &= \sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y) + 2 \sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X} \setminus \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y) \\ &\quad + \sum_{x \in \mathcal{X} \setminus \mathcal{X}_0} \sum_{y \in \mathcal{X} \setminus \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y) \\ &\leq \sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y) + 8\pi(\mathcal{X} \setminus \mathcal{X}_0). \end{aligned}$$

Using  $\pi(\mathcal{X}_0) \geq 1 - (\zeta/8)$ , we get

$$2 \left( \text{Var}_\pi(f) - \frac{\zeta}{2} \right) \leq \sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y).$$

Hence

$$\frac{\mathcal{E}(f, f)}{\text{Var}_\pi(f) - \frac{\zeta}{2}} \geq \frac{\sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) K(x, y)}{\sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) \pi(y)}.$$

For  $x \neq y$  in  $\mathcal{X}_0$ , let  $\gamma_{xy}$  be a canonical path linking  $x, y$  so that we can write, using the Cauchy-Schwarz inequality:

$$\begin{aligned} (f(y) - f(x))^2 &= \left( \sum_{e \in \gamma_{xy}} f(e_-) - f(e_+) \right)^2 \\ &\leq \sum_{e \in \gamma_{xy}} \frac{|\gamma_{xy}|}{\pi(e_-) K(e_-, e_+)} \times \pi(e_-) K(e_-, e_+) (f(e_-) - f(e_+))^2. \end{aligned}$$

It follows that

$$\begin{aligned} &\sum_{x \neq y} (f(y) - f(x))^2 \pi(x) \pi(y) \\ &\leq \sum_{x \neq y} \sum_{e \in \gamma_{xy}} \frac{\pi(x) \pi(y) |\gamma_{xy}|}{\pi(e_-) K(e_-, e_+)} \{ \pi(e_-) K(e_-, e_+) (f(e_-) - f(e_+))^2 \} \\ &= \sum_{e \in \mathcal{E}_0} \sum_{\gamma_{xy} \ni e} \left\{ \frac{\pi(x) \pi(y) |\gamma_{xy}|}{\pi(e_-) K(e_-, e_+)} \right\} \{ \pi(e_-) K(e_-, e_+) (f(e_-) - f(e_+))^2 \} \\ &\leq m(\mathcal{X}_0) \times \sum_{e \in \mathcal{E}_0} \pi(e_-) K(e_-, e_+) (f(e_-) - f(e_+))^2 \\ &\leq m(\mathcal{X}_0) \times \sum_{x \in \mathcal{X}_0} \sum_{y \in \mathcal{X}_0} (f(y) - f(x))^2 \pi(x) K(x, y). \end{aligned}$$

As a result we conclude that for all  $f : \mathcal{X} \rightarrow \mathbb{R}$ , with  $\|f\|_\infty \leq 1$ , and  $\text{Var}_\pi(f) > \zeta$

$$\frac{\mathcal{E}(f, f)}{\text{Var}_\pi(f) - \frac{\zeta}{2}} \geq \frac{1}{m(\mathcal{X}_0)}.$$

## REFERENCES

- ATCHADÉ, Y. and BHATTACHARYYA, A. (2019). An approach to large-scale quasi-bayesian inference with spike-and-slab priors.
- ATCHADÉ, Y. and WANG, L. (2023). A fast asynchronous Markov chain Monte Carlo sampler for sparse Bayesian inference. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **85** 1492–1516.
- ATCHADÉ, Y. F. (2021). Approximate spectral gaps for markov chain mixing times in high dimensions. *SIAM Journal on Mathematics of Data Science* **3** 854–872.
- BACH, F. (2010). Self-concordant analysis for logistic regression. *Electron. J. Statist.* **4** 384–414.
- BARBER, R. F., DRTON, M. and TAN, K. M. (2016). Laplace approximation in high-dimensional bayesian regression. In *Statistical Analysis for High-Dimensional Data: The Abel Symposium 2014*. Springer.
- BECK, A. and ELDAR, Y. C. (2013). Sparsity constrained nonlinear optimization: Optimality conditions and algorithms. *SIAM Journal on Optimization* **23** 1480–1509.
- BERTSIMAS, D., KING, A. and MAZUMDER, R. (2016). Best subset selection via a modern optimization lens. *The Annals of Statistics* **44** 813 – 852.
- BHADRA, A., DATTA, J., POLSON, N. G. and WILLARD, B. (2019). Lasso meets horseshoe: A survey. *Statistical Science* **34** pp. 405–427.  
URL <https://www.jstor.org/stable/26874188>
- BISWAS, N., JACOB, P. E. and VANETTI, P. (2019). Estimating convergence of markov chains with l-lag couplings. In *Advances in Neural Information Processing Systems* (H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox and R. Garnett, eds.), vol. 32. Curran Associates, Inc.
- BROOKS, S. P., GIUDICI, P. and ROBERTS, G. O. (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **65** 3–55.
- BROUSTE, A., DUTANG, C., MIENIEDOU, D. N. ET AL. (2021). Onestep: Le cam’s one-step estimation procedure. *R J.* **13** 366.
- CANDÈS, E. J. and PLAN, Y. (2009). Near-ideal model selection by L1 minimization. *The Annals of Statistics* **37** 2145 – 2177.
- CARLIN, B. P. and CHIB, S. (1995). Bayesian model choice via markov chain monte carlo methods. *Journal of the Royal Statistical Society. Series B (Methodological)* **57** 473–484.
- CHANG, H. and ZHOU, Q. (2024). Dimension-free relaxation times of informed mcmc samplers on discrete spaces.
- CHEN, J. and CHEN, Z. (2008). Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* **95** 759–771.
- CHEN, Y., DWIVEDI, R., WAINWRIGHT, M. J. and YU, B. (2020). Fast mixing of metropolized hamiltonian monte carlo: benefits of multi-step gradients. *J. Mach. Learn. Res.* **21**.

- DIACONIS, P. and STROOCK, D. (1991). Geometric bounds for eigenvalues of markov chains. *The Annals of Applied Probability* **1** 36–61.
- DOUC, R., MOULINES, E., PRIOURET, P. and SOULIER, P. (2018). *Markov chains*. Springer.
- FAN, J. and LV, J. (2010). A selective overview of variable selection in high dimensional feature space. *Statistica Sinica* 101–148.
- FOYGEL, R. and DRTON, M. (2015). High-dimensional ising model selection with bayesian information criteria. *Electronic Journal of Statistics* **9** 567–6–7.
- GEORGE, E. I. and MCCULLOCH, R. E. (1993). Variable selection via gibbs sampling. *Journal of the American Statistical Association* **88** 881–889.
- GREEN, P. J. (2003). Trans-dimensional markov chain monte carlo. *Oxford Statistical Science Series* 179–198.
- GRIFFIN, J. E., LATUSZYNSKI, K. G. and STEEL, M. F. J. (2020). In search of lost mixing time: adaptive Markov chain Monte Carlo schemes for Bayesian variable selection with very large  $p$ . *Biometrika* **108** 53–69.
- GUAN, Y. and STEPHENS, M. (2011). Bayesian variable selection regression for genome-wide association studies and other large-scale problems. *The Annals of Applied Statistics* **5** 1780 – 1815.
- GUO, J. (2018). *Some Contributions to High Dimensional Mixed Effects Logistic Regression Models*. Ph.D. thesis.
- HASTIE, T., TIBSHIRANI, R. and WAINWRIGHT, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC.
- HAZIMEH, H. and MAZUMDER, R. (2020). Fast best subset selection: Coordinate descent and local combinatorial optimization algorithms. *Oper. Res.* **68** 1517–1537.
- JACOB, P. E., O’LEARY, J. and ATCHADE, Y. F. (2020). Unbiased Markov Chain Monte Carlo Methods with Couplings. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **82** 543–600.
- JAVANMARD, A. and MONTANARI, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.* **15** 2869–2909.
- LAMNISOS, D., GRIFFIN, J. E. and STEEL, M. F. J. (2009). Transdimensional sampling algorithms for bayesian variable selection in classification problems with many more variables than observations. *Journal of Computational and Graphical Statistics* **18** 592–612.
- LAN, H., CHEN, M., FLOWERS, J. B., YANDELL, B. S., STAPLETON, D. S., MATA, C. M., MUI, E. T.-K., FLOWERS, M. T., SCHUELER, K. L., MANLY, K. F., WILLIAMS, R. W., KENDZIORSKI, C. and ATTIE, A. D. (2006). Combined expression trait correlations and expression quantitative trait locus mapping. *PLOS Genetics* **2** 1–11.  
URL <https://doi.org/10.1371/journal.pgen.0020006>
- LAWLER, G. and SOKAL, A. (1988). Bounds on the  $l_2$ spectrum for markov chains and markov processes: A generalization of cheeger’s inequality. *Transactions of the American Mathematical Society* **309** 557–580.

- LECAM, L. (1969). *Theorie Asymptotique De La Decision Statistique, Par Lucien M. Lecam*. Seminaire De Mathematiques Superieures, 33, Les presses de l'Université de Montreal.
- LIU, J. S. (1994). The collapsed gibbs sampler in bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association* **89** 958–966.
- LOVÁSZ, L. and SIMONOVITS, M. (1993). Random walks in a convex body and an improved volume algorithm. *Random Structures Algorithms* **4** 359–412.
- LUO, S. and CHEN, Z. (2011). Extended bic for linear regression models with diverging number of relevant features and high or ultra-high feature spaces. *Journal of Statistical Planning and Inference* **143**.
- MEINSHAUSEN, N. and YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37** 246–270.
- NARISSETY, N. N. and HE, X. (2014). Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics* **42** 789–817.  
URL <http://www.jstor.org/stable/43556304>
- NARISSETY, N. N., SHEN, J. and HE, X. (2018). Skinny gibbs: A consistent and scalable gibbs sampler for model selection. *Journal of the American Statistical Association* .
- O'HARA, R. B. and SILLANPÄÄ, M. J. (2009). A review of Bayesian variable selection methods: what, how and which. *Bayesian Analysis* **4** 85 – 117.
- POLSON, N. G. and SCOTT, J. G. (2011). 501Shrink Globally, Act Locally: Sparse Bayesian Regularization and Prediction. In *Bayesian Statistics 9*. Oxford University Press.
- RAY, K., SZABÓ, B. and CLARA, G. (2020). Spike and slab variational bayes for high dimensional logistic regression. *Advances in Neural Information Processing Systems* **33** 14423–14434.
- ROCKOVA, V. and GEORGE, E. I. (2018). The spike-and-slab lasso. *Journal of the American Statistical Association* **113** 431–444.
- SCHWARZ, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics* **6** 461 – 464.
- SINCLAIR, A. (1992). Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing* **1** 351–370.
- SINCLAIR, A. and JERRUM, M. (1989). Approximate counting, uniform generation and rapidly mixing Markov chains. *Inform. and Comput.* **82** 93–133.
- VAART, A. W. v. D. (1998). *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.
- VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42** 1166 – 1202.
- VAN DE VIJVER, M. J., HE, Y. D., VAN 'T VEER, L. J., DAI, H., HART, A. A., VOSKUIL, D. W., SCHREIBER, G. J., PETERSE, J. L., ROBERTS, C., MARTON, M. J., PARRISH, M., ATSMAN, D., WITTEVEEN, A., GLAS, A., DELAHAYE, L., VAN DER VELDE, T., BARTELINK, H., RODENHUIS, S., RUTGERS, E. T., FRIEND, S. H. and BERNARDS,

- R. (2002). A gene-expression signature as a predictor of survival in breast cancer. *New England Journal of Medicine* **347** 1999–2009. PMID: 12490681.  
URL <https://doi.org/10.1056/NEJMoa021967>
- VAN 'T VEER, D. H. V. D. V. M. E. A., L. (2002). Gene expression profiling predicts clinical outcome of breast cancer. *Nature* **415** 530–536.  
URL <https://doi.org/10.1038/415530a>
- VERSHYNIN, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.
- WAINWRIGHT, M. J. (2019). *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.
- WELCH, W. J. (1982). Algorithmic complexity: three np- hard problems in computational statistics. *Journal of Statistical Computation and Simulation* **15** 17–25.
- XIA, L., NAN, B. and LI, Y. (2020). A revisit to de-biased lasso for generalized linear models. *arXiv preprint arXiv:2006.12778* .
- YANG, Y., WAINWRIGHT, M. J. and JORDAN, M. I. (2016). On the computational complexity of high-dimensional bayesian variable selection. *Ann. Statist.* **44** 2497–2532.
- ZANELLA, G. and ROBERTS, G. (2019). Scalable Importance Tempering and Bayesian Variable Selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **81** 489–517.