Journal of the American Statistical Association

ISSN: 0162-1459 (Print) 1537-274X (Online) Journal homepage: http://www.tandfonline.com/loi/uasa20

Monte Carlo Simulation for Lasso-Type Problems by Estimator Augmentation

Qing Zhou

To cite this article: Qing Zhou (2014) Monte Carlo Simulation for Lasso-Type Problems by Estimator Augmentation, Journal of the American Statistical Association, 109:508, 1495-1516, DOI: 10.1080/01621459.2014.946035

To link to this article: http://dx.doi.org/10.1080/01621459.2014.946035

Accepted author version posted online: 05 Aug 2014.
Published online: 05 Aug 2014.

Submit your article to this journal

Article views: 773

View related articles

View Crossmark data

Citing articles: 1 View citing articles

Full Terms & Conditions of access and use can be found at
http://www.tandfonline.com/action/journalInformation?journalCode=uasa20

Download by: [UCLA Library] **Date:** 26 July 2017, At: 09:50

Monte Carlo Simulation for Lasso-Type Problems by Estimator Augmentation

Qing ZHOU

Regularized linear regression under the ℓ_1 penalty, such as the Lasso, has been shown to be effective in variable selection and sparse modeling. The sampling distribution of an ℓ_1-penalized estimator $\hat{\boldsymbol{\beta}}$ is hard to determine as the estimator is defined by an optimization problem that in general can only be solved numerically and many of its components may be exactly zero. Let $\mathbf{S}$ be the subgradient of the ℓ_1 norm of the coefficient vector $\boldsymbol{\beta}$ evaluated at $\hat{\boldsymbol{\beta}}$. We find that the joint sampling distribution of $\hat{\boldsymbol{\beta}}$ and $\mathbf{S}$, together called an augmented estimator, is much more tractable and has a closed-form density under a normal error distribution in both low-dimensional ($p \leq n$) and high-dimensional ($p > n$) settings. Given $\boldsymbol{\beta}$ and the error variance σ^2, one may employ standard Monte Carlo methods, such as Markov chain Monte Carlo (MCMC) and importance sampling (IS), to draw samples from the distribution of the augmented estimator and calculate expectations with respect to the sampling distribution of $\hat{\boldsymbol{\beta}}$. We develop a few concrete Monte Carlo algorithms and demonstrate with numerical examples that our approach may offer huge advantages and great flexibility in studying sampling distributions in ℓ_1-penalized linear regression. We also establish nonasymptotic bounds on the difference between the true sampling distribution of $\hat{\boldsymbol{\beta}}$ and its estimator obtained by plugging in estimated parameters, which justifies the validity of Monte Carlo simulation from an estimated sampling distribution even when $p \gg n \to \infty$.

KEY WORDS: Confidence interval; Importance sampling; Lasso; Markov chain Monte Carlo; *P*-value; Sampling distribution; Sparse linear model.

1. INTRODUCTION

Consider the linear regression model,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \varepsilon, \tag{1.1}$$

where $\mathbf{y}$ is an n-vector, $\mathbf{X}$ an $n \times p$ design matrix, $\boldsymbol{\beta} = (\beta_j)_{1:p}$ the vector of coefficients, and ε iid random errors with mean zero and variance σ^2. Recently, ℓ_1-penalized estimation methods (Tibshirani 1996; Chen, Donoho, and Saunders 1999) have been widely used to find sparse estimates of the coefficient vector. Given positive weights w_j, $j = 1, \ldots, p$, and a tuning parameter $\lambda > 0$, an ℓ_1-penalized estimator $\hat{\boldsymbol{\beta}} = (\hat{\beta}_j)_{1:p}$ is defined by minimizing the following penalized loss function,

$$\ell(\boldsymbol{\beta}) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + n\lambda \sum_{j=1}^{p} w_j|\beta_j|. \tag{1.2}$$

By different ways of choosing w_j, the estimator corresponds to the Lasso (Tibshirani 1996), the adaptive Lasso (Zou 2006), and the one-step linear local approximation (LLA) estimator (Zou and Li 2008) among others. We call such an estimator a Lasso-type estimator.

In many applications of ℓ_1-penalized regression, it is desired to quantify the uncertainty in the estimates. However, except for very special cases, the sampling distribution of a Lasso-type estimator is complicated and difficult to approximate. Closed-form approximations to the covariance matrices of the estimators in Tibshirani (1996), Fan and Li (2001), and Zou (2006) are unsatisfactory, as they all give zero variance for a zero component of the estimators and thus fail to quantify the uncertainty in variable selection. Theoretical results on finite-sample distributions and confidence sets of some Lasso-type estimators have been developed (Pötscher and Schneider 2009, 2010) but only under orthogonal designs, which clearly limits general applications of these results. The bootstrap can be used to approximate the sampling distribution of a Lasso-type estimator, in which numerical optimization is needed to minimize (1.2) for every bootstrap sample. Although there are efficient algorithms, such as the Lars (Efron et al. 2004), the homotopy algorithm (Osborne, Presnell, and Turlach 2000), and coordinate descent (Friedman et al. 2007; Wu and Lange 2008), to solve this optimization problem, it is still time-consuming to apply these algorithms hundreds or even thousands of times in bootstrap sampling. As pointed out by Knight and Fu (2000) and Chatterjee and Lahiri (2010), the bootstrap may not be consistent for estimating the sampling distribution of the Lasso under certain circumstances. To overcome this difficulty, a modified bootstrap (Chatterjee and Lahiri 2011) and a perturbation resampling approach (Minnier, Tian, and Cai 2011) have been proposed, both justified under a fixed-p asymptotic framework. Zhang and Zhang (2014) have developed methods for constructing confidence intervals for individual coefficients and their linear combinations in high-dimensional ($p > n$) regression with sufficient conditions for the asymptotic normality of the proposed estimators. There are several recent articles on significance test and confidence region construction for sparse high-dimensional linear models (Javanmard and Montanari 2013a, 2013b; van de Geer et al. 2013; Lockhart et al. 2014), all based on asymptotic distributions for various functions of the Lasso. On the other hand, knowledge on sampling distributions is also useful for distribution-based model selection with ℓ_1 penalization, as demonstrated by stability selection (Meinshausen and Bühlmann 2010) and the Bolasso (Bach 2008).

A possible alternative to the bootstrap or resampling is to simulate from a sampling distribution by Monte Carlo methods,

Qing Zhou is Associate Professor, Department of Statistics, University of California, Los Angeles, CA 90095, USA (E-mail: *zhou@stat.ucla.edu*).
This work was supported by NSF grants DMS-1055286 and DMS-1308376. The author thanks one referee for providing useful references.

Color versions of one or more of the figures in the article can be found online at *www.tandfonline.com/r/jasa*.

© 2014 American Statistical Association
Journal of the American Statistical Association
December 2014, Vol. 109, No. 508, Theory and Methods
DOI: 10.1080/01621459.2014.946035

such as Markov chain Monte Carlo (MCMC). An obvious obstacle to using these methods for a Lasso-type estimator is that its sampling distribution does not have a closed-form density. In this article, we study the joint distribution of a Lasso-type estimator $\hat{\boldsymbol{\beta}}$ and the subgradient $\mathbf{S}$ of $\|\boldsymbol{\beta}\|_1$ evaluated at $\hat{\boldsymbol{\beta}}$. Interestingly, this joint distribution has a density that can be calculated explicitly assuming a normal error distribution, regardless of the relative size between n and p. Thus, one can develop Monte Carlo algorithms to draw samples from this joint distribution and estimate various expectations of interest with respect to the sampling distribution of $\hat{\boldsymbol{\beta}}$, which is simply a marginal distribution. This approach offers great flexibility in studying the sampling distribution of a Lasso-type estimator. For instance, one may use importance sampling (IS) to accurately estimate a tail probability (small p-value) with respect to the sampling distribution under a null hypothesis, which can be orders of magnitude more efficient than any method directly targeting at the sampling distribution. Another potential advantage of this approach is that, at each iteration, an MCMC algorithm only evaluates a closed-form density, which is much faster than minimizing (1.2) numerically as used in the bootstrap. Furthermore, our method can be interpreted as an MCMC algorithm targeting at a multivariate normal distribution with locally reparameterized moves and hence is expected to be computationally tractable.

The remaining part of this article is organized as follows. After a high-level description of the basic idea, Section 2 derives the density of the joint distribution of $\hat{\boldsymbol{\beta}}$ and $\mathbf{S}$ in the low-dimensional setting with $p \le n$, and Section 3 develops MCMC algorithms for this setting. The density in the high-dimensional setting with $p > n$ is derived in Section 4. In Section 5, we construct applications of the high-dimensional result in p-value calculation for Lasso-type inference by IS. Numerical examples are provided in Sections 3 and 5 to demonstrate the efficiency of the Monte Carlo algorithms. Section 6 provides theoretical justifications for simulation from an estimated sampling distribution of the Lasso by establishing its consistency as $p \gg n \to \infty$. Section 7 includes generalizations to random designs, a connection to model selection consistency, and a Bayesian interpretation of the sampling distribution. The article concludes with a brief discussion and some remarks. Technical proofs are relegated to Section 8.

Notations for vectors and matrices are defined here. All vectors are regarded as column vectors. Let $A = \{j_1, \ldots, j_k\} \subseteq \{1, \ldots, m\}$ and $B = \{i_1, \ldots, i_\ell\} \subseteq \{1, \ldots, n\}$ be two index sets. For vectors $\mathbf{v} = (v_j)_{1:m}$ and $\mathbf{u} = (u_i)_{1:n}$, we define $\mathbf{v}_A = (v_j)_{j\in A} = (v_{j_1}, \ldots, v_{j_k})$, $\mathbf{v}_{-A} = (v_j)_{j\notin A}$, and $(\mathbf{v}_A, \mathbf{u}_B) = (v_{j_1}, \ldots, v_{j_k}, u_{i_1}, \ldots, u_{i_\ell})$. For a matrix $\mathbf{M} = (M_{ij})_{m\times n}$, write its columns as $\mathbf{M}_j$, $j = 1, \ldots, n$. Then, $\mathbf{M}_B = (\mathbf{M}_j)_{j\in B}$ extracts the columns in B, the submatrix $\mathbf{M}_{AB} = (M_{ij})_{i\in A, j\in B}$ extracts the rows in A and the columns in B, and $\mathbf{M}_{A\bullet} = (M_{ij})_{i\in A}$ extracts the rows in A. Furthermore, $\mathbf{M}_B^\mathsf{T}$ and $\mathbf{M}_{AB}^\mathsf{T}$ are understood as $(\mathbf{M}_B)^\mathsf{T}$ and $(\mathbf{M}_{AB})^\mathsf{T}$, respectively. We denote the row space, the null space, and the rank of $\mathbf{M}$ by row($\mathbf{M}$), null($\mathbf{M}$), and rank($\mathbf{M}$), respectively. Denote by diag($\mathbf{v}$) the $m \times m$ diagonal matrix with $\mathbf{v}$ as the diagonal elements, and by diag($\mathbf{M}$, $\mathbf{M}'$) the block diagonal matrix with $\mathbf{M}$ and $\mathbf{M}'$ as the diagonal blocks, where the submatrices $\mathbf{M}$ and $\mathbf{M}'$ may be of different sizes and may not be square. For a square matrix $\mathbf{M}$, diag($\mathbf{M}$) extracts the diagonal elements. Denote by $\mathbf{I}_n$ the $n \times n$ identity matrix.

2. ESTIMATOR AUGMENTATION

2.1 The Basic Idea

Let $\mathbf{W} = \text{diag}(w_1, \ldots, w_p)$. A minimizer $\hat{\boldsymbol{\beta}}$ of (1.2) is given by the Karush-Kuhn-Tucker (KKT) condition

$$\frac{1}{n}\mathbf{X}^\mathsf{T}\mathbf{y} = \frac{1}{n}\mathbf{X}^\mathsf{T}\mathbf{X}\hat{\boldsymbol{\beta}} + \lambda\mathbf{W}\mathbf{S}, \tag{2.1}$$

where $\mathbf{S} = (S_j)_{1:p}$ is the subgradient of the function $g(\boldsymbol{\beta}) = \|\boldsymbol{\beta}\|_1$ evaluated at the solution $\hat{\boldsymbol{\beta}}$. Therefore,

$$\begin{cases} S_j = \text{sgn}(\hat{\beta}_j) & \text{if } \hat{\beta}_j \neq 0, \\ S_j \in [-1, 1] & \text{if } \hat{\beta}_j = 0, \end{cases} \tag{2.2}$$

for $j = 1, \ldots, p$. Hereafter, we may simply call $\mathbf{S}$ the subgradient if the meaning is clear from context. Lemma 1 reviews a few basic facts about the uniqueness of $\hat{\boldsymbol{\beta}}$ and $\mathbf{S}$.

Lemma 1. For any $\mathbf{y}$, $\mathbf{X}$ and $\lambda > 0$, every minimizer $\hat{\boldsymbol{\beta}}$ of (1.2) gives the same fitted value $\mathbf{X}\hat{\boldsymbol{\beta}}$ and the same subgradient $\mathbf{S}$. Moreover, if the columns of $\mathbf{X}$ are in general position, then $\hat{\boldsymbol{\beta}}$ is unique for any $\mathbf{y}$ and $\lambda > 0$.

Proof. See Lemma 1 and Lemma 3 in Tibshirani (2013) for proof of the uniqueness of the fitted value $\mathbf{X}\hat{\boldsymbol{\beta}}$ and the uniqueness of $\hat{\boldsymbol{\beta}}$. Since $\mathbf{S}$ is a (vector-valued) function of $\mathbf{X}\hat{\boldsymbol{\beta}}$ from the KKT condition (2.1), it is also unique for fixed $\mathbf{y}$, $\mathbf{X}$ and λ. □

We regard $\hat{\boldsymbol{\beta}}$ and $\mathbf{S}$ together as the solution to Equation (2.1). Lemma 1 establishes that $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ is unique for any $\mathbf{y}$ assuming the columns of $\mathbf{X}$ are in general position (for a technical definition see Tibshirani 2013), regardless of the sizes of n and p. We call the vector $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ the augmented estimator in an ℓ_1-penalized regression problem. The augmented estimator will play a central role in our study of the sampling distribution of $\hat{\boldsymbol{\beta}}$.

Let $\mathbf{U} = \frac{1}{n}\mathbf{X}^\mathsf{T}\varepsilon = \frac{1}{n}\mathbf{X}^\mathsf{T}\mathbf{y} - \mathbf{C}\boldsymbol{\beta}$, where $\mathbf{C} = \frac{1}{n}\mathbf{X}^\mathsf{T}\mathbf{X}$ is the Gram matrix. By definition, $\mathbf{U} \in \text{row}(\mathbf{X})$. Rewrite the KKT condition as

$$\mathbf{U} = \mathbf{C}\hat{\boldsymbol{\beta}} + \lambda\mathbf{W}\mathbf{S} - \mathbf{C}\boldsymbol{\beta} \triangleq \mathbf{H}(\hat{\boldsymbol{\beta}}, \mathbf{S}; \boldsymbol{\beta}), \tag{2.3}$$

which shows that $\mathbf{U}$ is a function of $(\hat{\boldsymbol{\beta}}, \mathbf{S})$. On the other hand, $\mathbf{y}$ determines $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ only through $\mathbf{U}$, which implies that $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ is unique for any $\mathbf{U}$ as long as it is unique for any $\mathbf{y}$. Therefore, under the assumptions for the uniqueness of $\hat{\boldsymbol{\beta}}$, $\mathbf{H}$ is a bijection between $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ and $\mathbf{U}$. For a fixed $\mathbf{X}$, the only source of randomness in the linear model (1.1) is the noise vector ε, which determines the distribution of $\mathbf{U}$. With the bijection between $\mathbf{U}$ and $(\hat{\boldsymbol{\beta}}, \mathbf{S})$, one may derive the joint distribution of $(\hat{\boldsymbol{\beta}}, \mathbf{S})$, which has a closed-form density under a normal error distribution. Then, we develop Monte Carlo algorithms to sample from this joint distribution and obtain the sampling distribution of $\hat{\boldsymbol{\beta}}$. This is the key idea of this article, which works for both the low-dimensional setting ($p \le n$) and the high-dimensional setting ($p > n$). Although the basic strategy is the same, the technical details are slightly more complicated for the high-dimensional setting. For the sake of understanding, we first focus on the low-dimensional case in the remaining part of Section 2 and Section 3, and then generalize the results to the high-dimensional setting in Section 4.

Before going through all the technical details, we take a glimpse of the utility of this work in a couple concrete

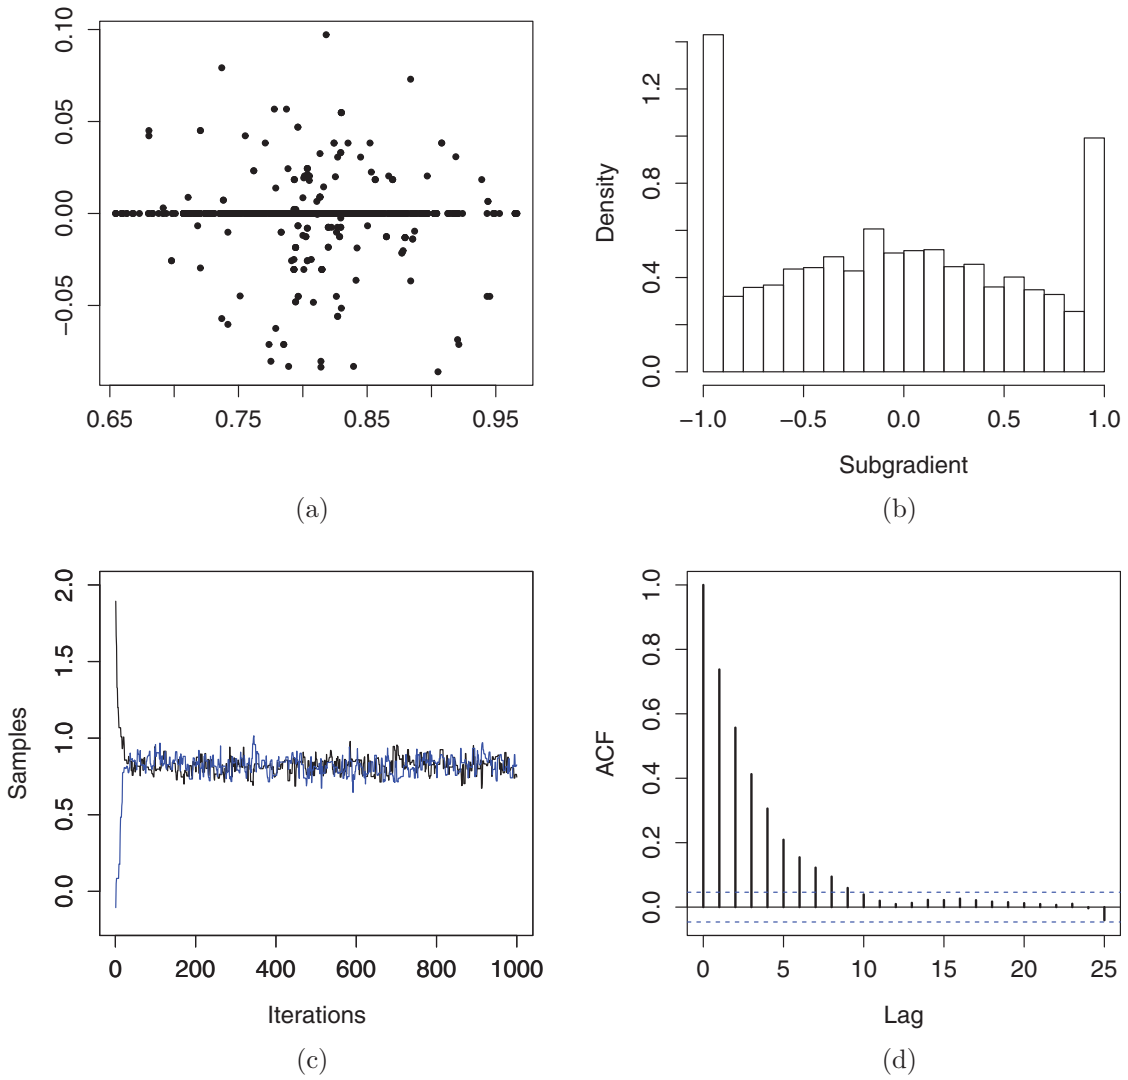

Figure 1. Demonstration of the Lasso sampler on a dataset with $p = 100$. (a) Scatterplot of the samples of $\hat{\beta}_1$ (x-axis) and $\hat{\beta}_{50}$ (y-axis); (b) histogram of the subgradient S_{50}; (c) two sample paths of $\hat{\beta}_1$ with diverse initial values; (d) a typical autocorrelation function.

examples. Given a design matrix $\mathbf{X}$ and a value of λ, our method gives a closed-form joint density π for the Lasso-type estimator $\hat{\boldsymbol{\beta}}$ and the subgradient $\mathbf{S}$ under assumed values of the true parameters (Theorems 1 and 2). Targeting at this density, we have developed MCMC algorithms, such as the Lasso sampler in Section 3.2, to draw samples from the joint distribution of $(\hat{\boldsymbol{\beta}}, \mathbf{S})$. Such MCMC samples allow for approximation of marginal distributions for a Lasso-type estimator. Figure 1 demonstrates the results of the Lasso sampler applied on a simulated dataset with $n = 500$, $p = 100$ and a normal error distribution. The scatterplot in Figure 1(a) confirms that $\hat{\beta}_j$ indeed may have a positive probability to be exactly zero in its sampling distribution. Accordingly, the distribution of the subgradient S_j in Figure 1(b) has a continuous density on $(-1, 1)$ and two point masses on ± 1. The fast mixing and low autocorrelation shown in the figure are surprisingly satisfactory for a simple MCMC algorithm in such a high-dimensional and complicated space ($\mathbb{R}^p \times 2^{\{1,\ldots,p\}}$, see (2.4)). Exploiting the explicit form of the bijection $\mathbf{H}$, we achieve the goal of sampling from the joint distribution of $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ via an MCMC algorithm essentially targeting at a multivariate normal distribution (Section 3.5). This approach does not need numerical optimization in any step, which makes it highly efficient compared to bootstrap or resampling-based methods. Another huge potential of our method is its ability to estimate tail probabilities, such as small p-values in a significance test (Section 5). Estimating tail probabilities is challenging for any simulation method. With a suitable proposal distribution, having an explicit density makes it possible to accurately estimate tail probabilities by importance weights. For example, our method can estimate a tail probability on the order of 10^{-20}, with a coefficient of variation around 2, by simulating only 5000 samples from a proposal distribution. This is absolutely impossible when bootstrapping the Lasso or simulating from the sampling distribution directly.

2.2 The Bijection

In the low-dimensional setting, we assume that $\operatorname{rank}(\mathbf{X}) = p \le n$, which guarantees that the columns of $\mathbf{X}$ are in general position.

Before writing down the bijection explicitly, we first examine the respective spaces for $\mathbf{U}$ and $(\hat{\boldsymbol{\beta}}, \mathbf{S})$. Under the assumption that $\operatorname{rank}(\mathbf{X}) = p$, the row space of $\mathbf{X}$ is simply $\mathbb{R}^p$, which is

the space for $\mathbf{U}$. Let $\mathcal{A} = \text{supp}(\hat{\boldsymbol{\beta}}) \triangleq \{j : \hat{\beta}_j \neq 0\}$ be the active set of $\hat{\boldsymbol{\beta}}$ and $\mathcal{I} = \{1, \ldots, p\} \setminus \mathcal{A}$ be the inactive set, that is, the set of the zero components of $\hat{\boldsymbol{\beta}}$. After removing the degeneracies among its components as given in (2.2), the vector $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ can be equivalently represented by the triple $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$. They are equivalent because from $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ one can unambiguously recover $(\hat{\boldsymbol{\beta}}, \mathbf{S})$, by setting $\hat{\boldsymbol{\beta}}_{\mathcal{I}} = \mathbf{0}$ and $\mathbf{S}_{\mathcal{A}} = \text{sgn}(\hat{\boldsymbol{\beta}}_{\mathcal{A}})$ (2.2), and vice versa. It is more convenient and transparent to work with this equivalent representation. One sees immediately that $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ lies in

$$\Omega = \{(\mathbf{b}_A, \mathbf{s}_I, A) : A \subseteq \{1, \ldots, p\}, \mathbf{b}_A \in (\mathbb{R} \setminus \{0\})^{|A|}, \mathbf{s}_I \in [-1, 1]^{p-|A|}\}, \tag{2.4}$$

where $I = \{1, \ldots, p\} \setminus A$. Hereafter, we always understand $(\mathbf{b}_A, \mathbf{s}_I, A)$ as the equivalent representation of $(\mathbf{b}, \mathbf{s}) = ((b_j)_{1:p}, (s_j)_{1:p})$ with $\text{supp}(\mathbf{b}) = A$ and $\mathbf{s}_A = \text{sgn}(\mathbf{b}_A)$. Clearly, $\Omega \subset \mathbb{R}^p \times 2^{\{1,\ldots,p\}}$, where $2^{\{1,\ldots,p\}}$ is the collection of all subsets of $\{1, \ldots, p\}$, and thus $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ lives in the product space of $\mathbb{R}^p$ and a finite discrete space.

Partition $\hat{\boldsymbol{\beta}}$ as $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \hat{\boldsymbol{\beta}}_{\mathcal{I}}) = (\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{0})$ and $\mathbf{S}$ as $(\mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}) = (\text{sgn}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}), \mathbf{S}_{\mathcal{I}})$. Then, the KKT condition (2.3) can be rewritten,

$$\mathbf{U} = (\mathbf{C}_{\mathcal{A}} \mid \mathbf{C}_{\mathcal{I}}) \begin{pmatrix} \hat{\boldsymbol{\beta}}_{\mathcal{A}} \\ \mathbf{0} \end{pmatrix} + \lambda (\mathbf{W}_{\mathcal{A}} \mid \mathbf{W}_{\mathcal{I}}) \begin{pmatrix} \mathbf{S}_{\mathcal{A}} \\ \mathbf{S}_{\mathcal{I}} \end{pmatrix} - \mathbf{C}\boldsymbol{\beta}, \tag{2.5}$$

$$= \mathbf{D}(\mathcal{A}) \begin{pmatrix} \hat{\boldsymbol{\beta}}_{\mathcal{A}} \\ \mathbf{S}_{\mathcal{I}} \end{pmatrix} + \lambda \mathbf{W}_{\mathcal{A}} \text{sgn}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}) - \mathbf{C}\boldsymbol{\beta} \triangleq \mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}), \tag{2.6}$$

where $\mathbf{D}(\mathcal{A}) = (\mathbf{C}_{\mathcal{A}} \mid \lambda \mathbf{W}_{\mathcal{I}})$ is a $p \times p$ matrix. Permuting the rows of $\mathbf{D}(\mathcal{A})$, one sees that

$$|\det \mathbf{D}(\mathcal{A})| = \det \begin{pmatrix} \mathbf{C}_{\mathcal{A}\mathcal{A}} & \mathbf{0} \\ \mathbf{C}_{\mathcal{I}\mathcal{A}} & \lambda \mathbf{W}_{\mathcal{I}\mathcal{I}} \end{pmatrix} = \lambda^{|\mathcal{I}|} \det(\mathbf{C}_{\mathcal{A}\mathcal{A}}) \prod_{j \in \mathcal{I}} w_j > 0 \tag{2.7}$$

if $\mathbf{C}_{\mathcal{A}\mathcal{A}} > 0$. Due to the equivalence between $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ and $(\hat{\boldsymbol{\beta}}, \mathbf{S})$, the map $\mathbf{H}$ defined here is essentially the same as the one defined in (2.3).

Lemma 2. If $\text{rank}(\mathbf{X}) = p$, then for any $\boldsymbol{\beta}$ and $\lambda > 0$, the mapping $\mathbf{H} : \Omega \to \mathbb{R}^p$ defined in (2.6) is a bijection that maps Ω onto $\mathbb{R}^p$.

Proof. For any $\mathbf{U} \in \mathbb{R}^p$, there is a unique solution $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ to Equation (2.3) if $\text{rank}(\mathbf{X}) = p$, and thus, a unique $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}) \in \Omega$ such that $\mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}) = \mathbf{U}$. For any $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}) \in \Omega$, $\mathbf{H}$ maps it into $\mathbb{R}^p$. □

It is helpful for understanding the map $\mathbf{H}$ to consider its inverse $\mathbf{H}^{-1}$ and its restriction to $\mathcal{A} = A$, where A is a fixed subset of $\{1, \ldots, p\}$. For any $\mathbf{U} \in \mathbb{R}^p$, if $\mathbf{H}^{-1}(\mathbf{U}; \boldsymbol{\beta}) = (\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$, then the unique solution to Equation (2.5) is $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$. Given a fixed A, $(\hat{\boldsymbol{\beta}}_A, \mathbf{S}_I)$ lives in the subspace

$$\Omega_A = \{(\mathbf{b}_A, \mathbf{s}_I) \in \mathbb{R}^p : \mathbf{b}_A \in (\mathbb{R} \setminus \{0\})^{|A|}, \mathbf{s}_I \in [-1, 1]^{p-|A|}\}. \tag{2.8}$$

Let $\mathbf{H}_A(\mathbf{b}_A, \mathbf{s}_I; \boldsymbol{\beta}) = \mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta})$ for $(\mathbf{b}_A, \mathbf{s}_I) \in \Omega_A$ and $U_A = \mathbf{H}_A(\Omega_A; \boldsymbol{\beta})$ be the image of Ω_A under the map $\mathbf{H}_A$. Now imagine we plug different $\mathbf{U} \in \mathbb{R}^p$ into Equation (2.5) and solve for $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$. Then, the set $\Omega_A \times \{A\}$ is the collection of all possible solutions such that $\text{supp}(\hat{\boldsymbol{\beta}}) = A$, the set U_A is the collection of all $\mathbf{U}$ that give these solutions, and $\mathbf{H}_A$ is a bijection between the two sets. It is easy to see that $\Omega = \bigcup_A \Omega_A \times \{A\}$, that is, $\{\Omega_A \times \{A\}\}$, for A extending over all subsets of $\{1, \ldots, p\}$, form a partition of the space Ω. The bijective nature of $\mathbf{H}$ implies that $\{U_A\}$ also form a partition of $\mathbb{R}^p$, the space of $\mathbf{U}$. Figure 2 illustrates the bijection $\mathbf{H}$ for $p = 2$ and the space partitioning by A. In this case, $\mathbf{H}_A$ map the four subspaces Ω_A for $A = \varnothing, \{1\}, \{2\}, \{1, 2\}$, each in a different $\mathbb{R}^2$, onto the space of $\mathbf{U}$ which is another $\mathbb{R}^2$.

Remark 1. The simple fact that $\mathbf{H}$ maps every point in Ω into $\text{row}(\mathbf{X}) = \mathbb{R}^p$ is crucial to the derivation of the sampling distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ in the low-dimensional setting. This means that every $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega$ is the solution to Equation (2.5) for $\mathbf{U} = \mathbf{u} = \mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta})$, and therefore one can simply find the probability density of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ at $(\mathbf{b}_A, \mathbf{s}_I, A)$ by the density of $\mathbf{U}$ at $\mathbf{u}$. This is not the case when $p > n$ (Section 4).

2.3 The Sampling Distribution

Now we can use the bijection $\mathbf{H}$ to find the distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ from the distribution of $\mathbf{U}$. Let ξ_k denote k-dimensional Lebesgue measure.

Theorem 1. Assume that $\text{rank}(\mathbf{X}) = p$ and let $f_{\mathbf{U}}$ be the probability density of $\mathbf{U}$ with respect to ξ_p. For $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega$, the joint distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ is given by

$$\begin{aligned} &P(\hat{\boldsymbol{\beta}}_A \in d\mathbf{b}_A, \mathbf{S}_I \in d\mathbf{s}_I, \mathcal{A} = A) \\ &\quad = f_{\mathbf{U}}(\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}))|\det \mathbf{D}(A)| \xi_p(d\mathbf{b}_A d\mathbf{s}_I) \\ &\quad \triangleq \pi(\mathbf{b}_A, \mathbf{s}_I, A) \xi_p(d\mathbf{b}_A d\mathbf{s}_I), \end{aligned} \tag{2.9}$$

and the distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathcal{A})$ is a marginal distribution given by

$$\begin{aligned} &P(\hat{\boldsymbol{\beta}}_A \in d\mathbf{b}_A, \mathcal{A} = A) \\ &\quad = \left[\int_{[-1,1]^{p-|A|}} \pi(\mathbf{b}_A, \mathbf{s}_I, A) \xi_{p-|A|}(d\mathbf{s}_I) \right] \xi_{|A|}(d\mathbf{b}_A). \end{aligned} \tag{2.10}$$

Proof. Let $\mathbf{u} = \mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}) = \mathbf{H}_A(\mathbf{b}_A, \mathbf{s}_I; \boldsymbol{\beta})$. From (2.6) and (2.8), one sees that for any fixed A, $b_j \neq 0$ for all $j \in A$ and $\mathbf{H}_A$ is differentiable. Differentiating $\mathbf{u}$ with respect to $(\mathbf{b}_A, \mathbf{s}_I)$,

$$d\mathbf{u} = \frac{\partial \mathbf{H}_A}{\partial (\mathbf{b}_A, \mathbf{s}_I)^{\mathsf{T}}} \begin{pmatrix} d\mathbf{b}_A \\ d\mathbf{s}_I \end{pmatrix} = \mathbf{D}(A) \begin{pmatrix} d\mathbf{b}_A \\ d\mathbf{s}_I \end{pmatrix}$$

and thus $\xi_p(d\mathbf{u}) = |\det \mathbf{D}(A)| \xi_p(d\mathbf{b}_A d\mathbf{s}_I)$. Since $\mathbf{H}$ and $\mathbf{H}_A : \Omega_A \to U_A$ are bijections, a change of variable gives

$$\begin{aligned} P(\hat{\boldsymbol{\beta}}_A \in d\mathbf{b}_A, \mathbf{S}_I \in d\mathbf{s}_I, \mathcal{A} = A) &= P(\mathbf{U} \in d\mathbf{u}) \\ &= f_{\mathbf{U}}(\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}))|\det \mathbf{D}(A)| \xi_p(d\mathbf{b}_A d\mathbf{s}_I). \end{aligned}$$

Integrating (2.9) over $\mathbf{s}_I \in [-1, 1]^{p-|A|}$ gives (2.10). □

Remark 2. Equation (2.9) gives the joint distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ and effectively the joint distribution of $(\hat{\boldsymbol{\beta}}, \mathbf{S})$. The density $\pi(\mathbf{b}_A, \mathbf{s}_I, A)$ is defined with respect to the product of ξ_p and counting measure on $2^{\{1,\ldots,p\}}$. Analogously, the sampling distribution of $\hat{\boldsymbol{\beta}}$ is given by the distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathcal{A})$ in (2.10). To be rigorous, (2.9) is derived by assuming that $(\mathbf{b}_A, \mathbf{s}_I)$ is an interior point of Ω_A. Note that $(\mathbf{b}_A, \mathbf{s}_I) \in \Omega_A$ is not in the interior if and only if $|s_j| = 1$ for some $j \in I$, and thus the Lebesgue

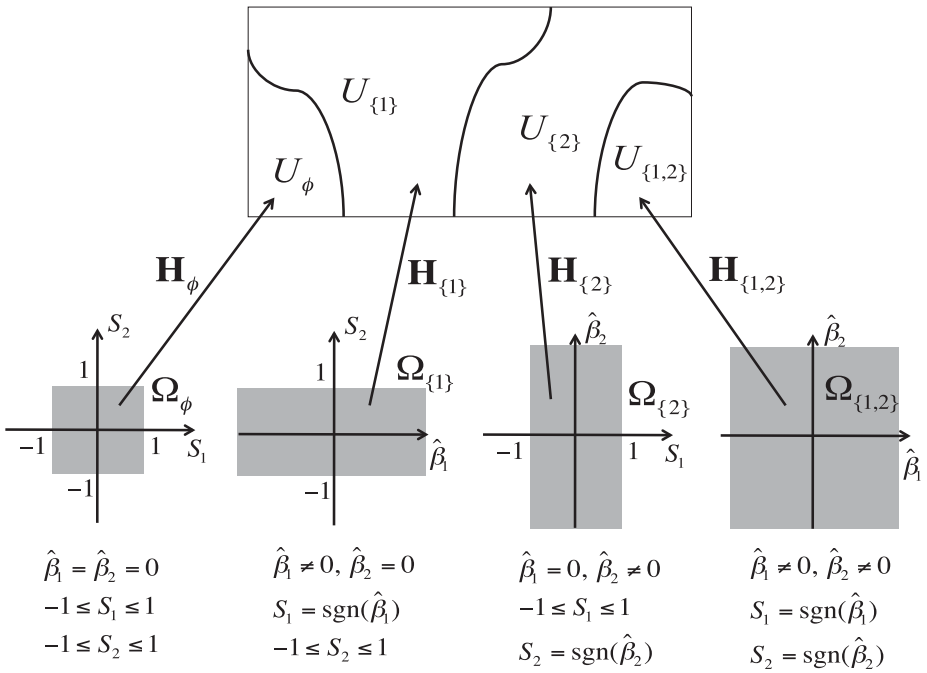

Figure 2. The bijection $\mathbf{H}$, its restrictions $\mathbf{H}_A$, the four subspaces Ω_A (shaded areas) and the corresponding partition in the space of $\mathbf{U}$ for $p = 2$.

measure of the set of these points is zero. Therefore, it will cause no problem at all to use π as the density for all points in Ω. The joint distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ has at least two nice properties which make it much more tractable than the distribution of $\hat{\boldsymbol{\beta}}$. First, the density π does not involve multidimensional integral and has a closed-form expression that can be calculated explicitly if $f_{\mathbf{U}}$ is given. Second, the continuous components $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}})$ always have the same dimension $(= p)$ for any value of $\mathcal{A}$, while $\hat{\boldsymbol{\beta}}_{\mathcal{A}}$ lives in $\mathbb{R}^{|\mathcal{A}|}$ whose dimension changes with $\mathcal{A}$. These two properties are critical to the development of MCMC to sample from π. See Section 3.1 for more discussion. We explicitly include the dominating Lebesgue measure to clarify the dimension of a density.

Remark 3. The distribution of $\hat{\boldsymbol{\beta}} = (\hat\beta_1, \ldots, \hat\beta_p)$ in (2.10) is essentially defined for each A. In many problems, one may be interested in the marginal distribution of $\hat\beta_j$ such as for calculating p-values and constructing confidence intervals. To obtain such a marginal distribution, we need to sum over all possible active sets, which cannot be done analytically. Our strategy is to draw samples from the joint distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ by a Monte Carlo method. Then from the Monte Carlo samples one can easily approximate any marginal distribution of interest, such as that of $\hat\beta_j$. This is exactly our motivation for estimator augmentation, which is in spirit similar to the use of auxiliary variables in the MCMC literature.

To further help our understanding of the density π, consider a few conditional and marginal distributions derived from the joint distribution (2.9). First, the sampling distribution of the active set $\mathcal{A}$ is given by

$$P(\mathcal{A} = A) = \int_{\Omega_A} \pi(\mathbf{b}_A, \mathbf{s}_I, A)\xi_p(d\mathbf{b}_A d\mathbf{s}_I) \triangleq Z_A, \qquad (2.11)$$

where Ω_A is the subspace for $(\hat{\boldsymbol{\beta}}_A, \mathbf{S}_I)$ defined in (2.8). In other words, Z_A is the probability of $\Omega_A \times \{A\}$ with respect to the joint distribution π. Second, the conditional density of $(\hat{\boldsymbol{\beta}}_A, \mathbf{S}_I)$ given $\mathcal{A} = A$ (with respect to ξ_p) is

$$\pi(\mathbf{b}_A, \mathbf{s}_I \mid A) = \frac{1}{Z_A}\pi(\mathbf{b}_A, \mathbf{s}_I, A) \propto f_{\mathbf{U}}(\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta})) \qquad (2.12)$$

for $(\mathbf{b}_A, \mathbf{s}_I) \in \Omega_A \subset \mathbb{R}^p$. Using $p = 2$ as an illustration, the joint density π is defined over all four shaded areas in Figure 2, while a conditional density $\pi(\cdot \mid A)$ is defined on each one of them. To give a concrete probability calculation, for $a_2 > a_1 > 0$,

$$\begin{aligned} P(\hat\beta_1 \in [a_1, a_2], \hat\beta_2 = 0) &= P(\hat\beta_1 \in [a_1, a_2], \mathcal{A} = \{1\}) \\ &= \int_{a_1}^{a_2}\int_{-1}^{1} \pi(b_1, s_2, \{1\}) ds_2 db_1, \end{aligned}$$

which is an integral over the rectangle $[a_1, a_2] \times [-1, 1]$ in $\Omega_{\{1\}}$ (Figure 2). Clearly, this probability can be approximated by Monte Carlo integration if we have enough samples from π.

Remark 4. We emphasize that the weights w_j and the tuning parameter λ are assumed to be fixed in Theorem 1. For the adaptive Lasso, one may choose $w_j = |\tilde\beta_j|^{-1}$ based on an initial estimate $\tilde\beta_j$. The distribution (2.9) is valid for the adaptive Lasso only if we ignore the randomness in $\tilde\beta_j$ and regard w_j as constants during the repeated sampling procedure.

2.4 Normal Errors

Denote by $\mathcal{N}_k(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ the k-variate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, and by $\phi_k(\mathbf{z}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ its probability density function. If the error $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$ and $\mathrm{rank}(\mathbf{X}) = p$, then $\mathbf{U} \sim \mathcal{N}_p(\mathbf{0}, \frac{\sigma^2}{n}\mathbf{C})$. In this case, the joint density π (2.9) has a closed-form expression. Recall that $\mathbf{s}_A =$

$\operatorname{sgn}(\mathbf{b}_A)$ and define

$$\boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta}) = [\mathbf{D}(A)]^{-1}(\mathbf{C}\boldsymbol{\beta} - \lambda \mathbf{W}_A \mathbf{s}_A), \tag{2.13}$$

$$\boldsymbol{\Sigma}(A; \sigma^2) = \frac{\sigma^2}{n}[\mathbf{D}(A)]^{-1}\mathbf{C}[\mathbf{D}(A)]^{-\mathsf{T}}. \tag{2.14}$$

Corollary 1. If $\operatorname{rank}(\mathbf{X}) = p$ and $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, then the joint density of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ is

$$\pi(\mathbf{b}_A, \mathbf{s}_I, A) = \phi_p(\mathbf{z}; \boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta}), \boldsymbol{\Sigma}(A; \sigma^2))\mathbf{1}((\mathbf{z}, A) \in \Omega), \tag{2.15}$$

where $\mathbf{z} = (\mathbf{b}_A, \mathbf{s}_I) \in \mathbb{R}^p$ and $\mathbf{1}(\cdot)$ is an indicator function.

Proof. First note that

$$\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}) = \mathbf{D}(A)[\mathbf{z} - \boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta})]. \tag{2.16}$$

Under the assumptions, $\mathbf{U} \sim \mathcal{N}_p(\mathbf{0}, \frac{\sigma^2}{n}\mathbf{C})$. By Theorem 1,

$$\begin{aligned}
&\pi(\mathbf{b}_A, \mathbf{s}_I, A)\\
&\quad = \phi_p\left(\mathbf{D}(A)[\mathbf{z} - \boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta})]; \mathbf{0}, n^{-1}\sigma^2\mathbf{C}\right) |\det \mathbf{D}(A)|\\
&\quad = \phi_p\left(\mathbf{z}; \boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta}), n^{-1}\sigma^2[\mathbf{D}(A)]^{-1}\mathbf{C}[\mathbf{D}(A)]^{-\mathsf{T}}\right)\\
&\quad = \phi_p(\mathbf{z}; \boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta}), \boldsymbol{\Sigma}(A; \sigma^2))
\end{aligned}$$

for $(\mathbf{b}_A, \mathbf{s}_I, A) = (\mathbf{z}, A) \in \Omega$. □

Without the normal error assumption, Corollary 1 is still a good approximation when n is large and p is fixed, since $\sqrt{n}\mathbf{U} \xrightarrow{d} \mathcal{N}_p(\mathbf{0}, \sigma^2\mathbf{C})$ assuming $\frac{1}{n}\mathbf{X}^{\mathsf{T}}\mathbf{X} \to \mathbf{C} > 0$ as $n \to \infty$.

Note that both the continuous components $\mathbf{z}$ and the active set A are arguments of the density (2.15). For different A and $\mathbf{s}_A$, the normal density ϕ_p has different parameters. Given a particular A^* and $\mathbf{s}^* \in \{\pm 1\}^{|A^*|}$, let $I^* = \{1, \ldots, p\} \setminus A^*$ and

$$\Omega_{A^*,\mathbf{s}^*} = \{(\mathbf{b}_{A^*}, \mathbf{s}_{I^*}) \in \Omega_{A^*} : \operatorname{sgn}(\mathbf{b}_{A^*}) = \mathbf{s}^*\}. \tag{2.17}$$

Then $\Omega_{A^*,\mathbf{s}^*} \times \{A^*\}$ is the subset of Ω corresponding to the event $\{\mathcal{A} = A^*, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_{A^*}) = \mathbf{s}^*\}$. For $\mathbf{z} \in \Omega_{A^*,\mathbf{s}^*}$, the density $\pi(\mathbf{z}, A^*)$ is identical to $\phi_p(\mathbf{z}; \boldsymbol{\mu}(A^*, \mathbf{s}^*; \boldsymbol{\beta}), \boldsymbol{\Sigma}(A^*; \sigma^2))$, that is,

$$\begin{aligned}
\pi(\mathbf{z}, A^*)\mathbf{1}(\mathbf{z} \in \Omega_{A^*,\mathbf{s}^*}) = \phi_p(\mathbf{z}; \boldsymbol{\mu}(A^*, \mathbf{s}^*; \boldsymbol{\beta}),\\
\boldsymbol{\Sigma}(A^*; \sigma^2))\mathbf{1}(\mathbf{z} \in \Omega_{A^*,\mathbf{s}^*}).
\end{aligned}$$

Intuitively, this is because $\mathbf{H}$ restricted to $A = A^*$ and $\mathbf{s}_{A^*} = \mathbf{s}^*$ is simply an affine map [see (2.16)]. Consequently, the probability of $\Omega_{A^*,\mathbf{s}^*} \times \{A^*\}$ with respect to π is

$$\begin{aligned}
&P(\mathcal{A} = A^*, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_{A^*}) = \mathbf{s}^*)\\
&\quad = \int_{\Omega_{A^*,\mathbf{s}^*}} \phi_p(\mathbf{z}; \boldsymbol{\mu}(A^*, \mathbf{s}^*; \boldsymbol{\beta}), \boldsymbol{\Sigma}(A^*; \sigma^2))\xi_p(d\mathbf{z}),
\end{aligned} \tag{2.18}$$

and $[\hat{\boldsymbol{\beta}}_{A^*}, \mathbf{S}_{I^*} \mid \mathcal{A} = A^*, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_{A^*}) = \mathbf{s}^*]$ is the truncated $\mathcal{N}_p(\boldsymbol{\mu}(A^*, \mathbf{s}^*; \boldsymbol{\beta}), \boldsymbol{\Sigma}(A^*; \sigma^2))$ on $\Omega_{A^*,\mathbf{s}^*}$. For $p = 2$, if $A^* = \{1\}$ and $\mathbf{s}^* = -1$, the region $\Omega_{\{1\},-1} = (-\infty, 0) \times [-1, 1]$ is the left half of the $\Omega_{\{1\}}$ in Figure 2 and the density π restricted to this region is the same as the part of a bivariate normal density on the same region.

If $\mathbf{C} = \mathbf{I}_p$ and $\mathbf{W} = \mathbf{I}_p$, the Lasso is equivalent to soft-thresholding the ordinary least-squares estimator $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}} = (\hat{\beta}_j^{\mathrm{OLS}})_{1:p}$. In this case, (2.13) and (2.14) have simpler forms:

$$\boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta}) = \begin{pmatrix} \boldsymbol{\beta}_A - \lambda \mathbf{s}_A \\ \lambda^{-1}\boldsymbol{\beta}_I \end{pmatrix},$$

$$\boldsymbol{\Sigma}(A; \sigma^2) = \frac{\sigma^2}{n}\begin{pmatrix} \mathbf{I}_{|A|} & \mathbf{0} \\ \mathbf{0} & \lambda^{-2}\mathbf{I}_{|I|} \end{pmatrix}.$$

By (2.18) we find, for example,

$$\begin{aligned}
&P(\mathcal{A} = A, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_A) = (1, \ldots, 1))\\
&\quad = \prod_{j \in A} \int_0^\infty \phi\left(b_j; \beta_j - \lambda, \frac{\sigma^2}{n}\right) db_j\\
&\qquad \cdot \prod_{j \in I} \int_{-1}^1 \phi\left(s_j; \frac{\beta_j}{\lambda}, \frac{\sigma^2}{\lambda^2 n}\right) ds_j\\
&\quad = \prod_{j \in A} P(\hat{\beta}_j^{\mathrm{OLS}} > \lambda) \cdot \prod_{j \in I} P\left(|\hat{\beta}_j^{\mathrm{OLS}}| \le \lambda\right),
\end{aligned}$$

where the last equality is due to that $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}} \sim \mathcal{N}_p(\boldsymbol{\beta}, n^{-1}\sigma^2\mathbf{I}_p)$. One sees that our result is consistent with that obtained directly from soft-thresholding each component of $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ by λ.

2.5 Estimation

To apply Theorem 1 in practice, one needs to estimate $f_{\mathbf{U}}$ and $\boldsymbol{\beta}$ if they are not given. Suppose that $f_{\mathbf{U}}$ is estimated by $\hat{f}_{\mathbf{U}}$ and $\boldsymbol{\beta}$ is estimated by $\check{\boldsymbol{\beta}}$. Then, the corresponding estimate of the density π is

$$\hat{\pi}(\mathbf{b}_A, \mathbf{s}_I, A) = \hat{f}_{\mathbf{U}}(\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \check{\boldsymbol{\beta}}))|\det \mathbf{D}(A)|. \tag{2.19}$$

Since $\mathbb{E}(\mathbf{U}) = \mathbf{0}$ and $\operatorname{Var}(\sqrt{n}\mathbf{U}) = \sigma^2\mathbf{C}$, estimating $f_{\mathbf{U}}$ reduces to estimating σ^2 when ε is normally distributed or when the sample size n is large. A consistent estimator of σ^2 can be constructed given a consistent estimator of $\boldsymbol{\beta}$. For example, when $p < n$ one may use

$$\hat{\sigma}^2 = \frac{\|\mathbf{y} - \mathbf{X}\check{\boldsymbol{\beta}}\|_2^2}{n - p}, \tag{2.20}$$

provided that $\check{\boldsymbol{\beta}}$ is consistent for $\boldsymbol{\beta}$. If ε does not follow a normal distribution, one can apply other parametric or nonparametric methods to estimate $f_{\mathbf{U}}$. Here, we propose a bootstrap-based approach under the assumption that $\mathbf{U}$ is elliptically symmetric. That is, $\tilde{\mathbf{U}} = \mathbf{C}^{-1/2}\mathbf{U}$ is spherically symmetric: For $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^p$, if $\|\mathbf{v}_1\|_2 = \|\mathbf{v}_2\|_2$ then $f_{\tilde{\mathbf{U}}}(\mathbf{v}_1) = f_{\tilde{\mathbf{U}}}(\mathbf{v}_2)$, where $f_{\tilde{\mathbf{U}}}$ is the density of $\tilde{\mathbf{U}}$. Generate bootstrap samples, $\varepsilon^{(i)} = (\varepsilon_1^{(i)}, \ldots, \varepsilon_n^{(i)})$ for $i = 1, \ldots, K$, by resampling with replacement from $\hat{\varepsilon} = (\hat{\varepsilon}_1, \ldots, \hat{\varepsilon}_n) = \mathbf{y} - \mathbf{X}\check{\boldsymbol{\beta}}$, and calculate $\tilde{\mathbf{U}}^{(i)} = \frac{1}{n}\mathbf{C}^{-1/2}\mathbf{X}^{\mathsf{T}}\varepsilon^{(i)}$ for each i. Given $0 = h_0 < h_1 < \cdots < h_M < \infty$, let $K_m = |\{i : h_{m-1} \le \|\tilde{\mathbf{U}}^{(i)}\|_2 < h_m\}|$ for $m = 1, \ldots, M$. The density of $\tilde{\mathbf{U}}$ is then estimated by

$$\hat{f}_{\tilde{\mathbf{U}}}(\mathbf{v}) \propto \sum_{m=1}^M \frac{K_m}{h_m^p - h_{m-1}^p}\mathbf{1}(h_{m-1} \le \|\mathbf{v}\|_2 < h_m) \tag{2.21}$$

for $\|\mathbf{v}\|_2 \in [0, h_M)$. The density for $\|\mathbf{v}\|_2 \ge h_M$ can be estimated by linear extrapolation of $\log \hat{f}_{\tilde{\mathbf{U}}}$. Finally, set $\hat{f}_{\mathbf{U}}(\mathbf{u}) = \hat{f}_{\tilde{\mathbf{U}}}(\mathbf{C}^{-1/2}\mathbf{u})(\det \mathbf{C})^{-1/2}$.

In general, estimating $f_{\mathbf{U}}$ is difficult when p is large. One may have to assume some parametric density for $\mathbf{U}$, which reduces

the problem to the estimation of a few unknown parameters. Besides normality, one may assume that $\mathbf{U}$ follows a multivariate t distribution, which is motivated from a Bayesian perspective to be discussed in Section 7.3.

Sampling from π or $\hat{\pi}$ can be very useful for statistical inference based on a Lasso-type estimator. We may directly draw $(\boldsymbol{\beta}^*, \mathbf{S}^*)$ from $\hat{\pi}$ given $(\check{\boldsymbol{\beta}}, \hat{\sigma})$ and use the conditional distribution $[(\boldsymbol{\beta}^* - \check{\boldsymbol{\beta}}) \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$ to construct confidence regions around $\hat{\boldsymbol{\beta}}$. Under some assumptions, the conditional distribution $[(\boldsymbol{\beta}^* - \check{\boldsymbol{\beta}}) \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$ provides a valid approximation to the true sampling distribution of $(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$. We derive nonasymptotic error bounds for this approximation in Section 6 after the development of our method in the high-dimensional setting. If $\boldsymbol{\beta}$ is specified in the null hypothesis in a significance test, then samples from π can be used to calculate p-values. This aspect will be explored in Section 5.

3. MCMC ALGORITHMS

In this section, we develop MCMC algorithms to sample from π given $\boldsymbol{\beta}$ and $f_{\mathbf{U}}$ (or σ^2). Before that, we first introduce a direct sampling approach which includes the residual bootstrap method as a special case.

Routine 1 (Direct sampler). Assume the error distribution is $\mathcal{D}(\mathbf{0}, \sigma^2\mathbf{I}_n)$. For $t = 1, \ldots, L$

1. draw $\varepsilon^{(t)} \sim \mathcal{D}(\mathbf{0}, \sigma^2\mathbf{I}_n)$ and set $\mathbf{y}^{(t)} = \mathbf{X}\boldsymbol{\beta} + \varepsilon^{(t)}$;
2. find the minimizer $\hat{\boldsymbol{\beta}}^{(t)}$ of (1.2) with $\mathbf{y}^{(t)}$ in place of $\mathbf{y}$;
3. if needed, calculate the subgradient vector $\mathbf{S}^{(t)} = (n\lambda\mathbf{W})^{-1}\mathbf{X}^{\mathsf{T}}(\mathbf{y}^{(t)} - \mathbf{X}\hat{\boldsymbol{\beta}}^{(t)})$.

This approach directly draws $\mathbf{y}^{(t)}$ from its sampling distribution and requires a numerical optimization algorithm in Step 2 for each sample. Moreover, Step 1 will be complicated if we cannot draw independent samples from $\mathcal{D}(\mathbf{0}, \sigma^2\mathbf{I}_n)$. If $\varepsilon^{(t)}$ is drawn by resampling residuals, then Routine 1 is equivalent to the bootstrap method of Knight and Fu (2000).

As the density $\pi(\mathbf{b}_A, \mathbf{s}_I, A)$ (2.9) has a closed-form expression given $\boldsymbol{\beta}$ and $f_{\mathbf{U}}$, MCMC and IS can be applied to sample from and calculate expectations with respect to the distribution. These methods may offer much more flexible and efficient alternatives to the direct sampling approach, although the samples are either dependent or weighted. In what follows, we propose a few special designs targeting at different applications to exemplify the use of MCMC methods. Examples of IS will be given in Section 5 under the high-dimensional setting.

3.1 Reversibility

Our goal is to design a reversible Markov chain on the space Ω, which is composed of a finite number of subspaces Ω_A, each having the same dimension p. Therefore, moves with an ordinary Metropolis-Hastings (MH) ratio are sufficient, which can be seen as follows. For any $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega$, let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_p)$ with components given by

$$\theta_j = \begin{cases} b_j & \text{if } j \in A \\ s_j & \text{otherwise,} \end{cases} \tag{3.1}$$

that is, $\boldsymbol{\theta}_A = \mathbf{b}_A$ and $\boldsymbol{\theta}_I = \mathbf{s}_I$. Then, our target distribution is $\pi(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A)\xi_p(d\boldsymbol{\theta})$. Suppose that $(\boldsymbol{\theta}, A)$ is the current state and we have a proposal for a new state $(\boldsymbol{\theta}^\dagger, A^\dagger)$. In general, the proposal may only change some components of $\boldsymbol{\theta}$, say θ_j for $j \in B \subseteq \{1, \ldots, p\}$, such that $\boldsymbol{\theta}^\dagger_{-B} = \boldsymbol{\theta}_{-B}$. Let $q((\boldsymbol{\theta}, A), (\boldsymbol{\theta}^\dagger, A^\dagger))$ be the density of this proposal with respect to $\xi_{|B|}$ and $I^\dagger = \{1, \ldots, p\} \setminus A^\dagger$. The MH ratio in terms of probability measures is

$$\begin{aligned} &\min\left\{1, \frac{\pi(\boldsymbol{\theta}^\dagger_{A^\dagger}, \boldsymbol{\theta}^\dagger_{I^\dagger}, A^\dagger)\xi_p(d\boldsymbol{\theta}^\dagger)}{\pi(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A)\xi_p(d\boldsymbol{\theta})} \frac{q((\boldsymbol{\theta}^\dagger, A^\dagger), (\boldsymbol{\theta}, A))\xi_{|B|}(d\boldsymbol{\theta}_B)}{q((\boldsymbol{\theta}, A), (\boldsymbol{\theta}^\dagger, A^\dagger))\xi_{|B|}(d\boldsymbol{\theta}^\dagger_B)}\right\} \\ &= \min\left\{1, \frac{\pi(\boldsymbol{\theta}^\dagger_{A^\dagger}, \boldsymbol{\theta}^\dagger_{I^\dagger}, A^\dagger)}{\pi(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A)} \frac{q((\boldsymbol{\theta}^\dagger, A^\dagger), (\boldsymbol{\theta}, A))}{q((\boldsymbol{\theta}, A), (\boldsymbol{\theta}^\dagger, A^\dagger))} \frac{\xi_{p-|B|}(d\boldsymbol{\theta}^\dagger_{-B})}{\xi_{p-|B|}(d\boldsymbol{\theta}_{-B})}\right\}. \end{aligned} \tag{3.2}$$

As $\boldsymbol{\theta}^\dagger_{-B} = \boldsymbol{\theta}_{-B}$, the dominating measures in (3.2) cancel out and the ratio reduces to a standard MH ratio involving only densities.

Now we see that our strategy of estimator augmentation plays two roles in MCMC sampling. First, $\mathbf{S}_{\mathcal{I}}$ plays the role of an auxiliary variable: The target distribution π for $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ has a closed-form density which allows one to design an MCMC algorithm, while the distribution of interest, that for $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathcal{A})$, is a marginal distribution of π without a closed-form density. Second, $\mathbf{S}_{\mathcal{I}}$ also plays the role of dimension matching so that the continuous components $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}})$ always have the same dimension in any subspace. This eliminates the need for reversible jump MCMC (Green 1995). On the contrary, if we were to sample $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathcal{A})$ (assuming a closed-form approximation to its density), moves between two subspaces of different dimensions would require reversible jumps, which are usually much harder to design.

3.2 The MH Lasso Sampler

We develop an MH algorithm, called the MH Lasso sampler (MLS), with coordinate-wise update. That is, to sequentially update each θ_j, $j = 1, \ldots, p$, while holding other components fixed. Suppose the current state is $(\boldsymbol{\theta}, A)$. We design four moves to propose a new state $(\boldsymbol{\theta}^\dagger, A^\dagger)$, which are grouped into two types according to whether $A^\dagger = A$ or not. In the following proposals, $\theta^\dagger_j = b^\dagger_j$ if $j \in A^\dagger$ and $\theta^\dagger_j = s^\dagger_j$ otherwise.

Definition 1. Proposals in the MLS for a given $j \in \{1, \ldots, p\}$.

- Parameter-update proposals: (P1) If $j \in A$, draw $b^\dagger_j \sim \mathcal{N}(b_j, \tau_j^2)$. (P2) If $j \notin A$, draw $s^\dagger_j \sim \text{Unif}(-1, 1)$. Set $A^\dagger = A$ in both (P1) and (P2).
- Model-update proposals: (P3) If $j \in A$, set $A^\dagger = A \setminus \{j\}$ and draw $s^\dagger_j \sim \text{Unif}(-1, 1)$. (P4) If $j \notin A$, set $A^\dagger = A \cup \{j\}$ and draw $b^\dagger_j \sim \mathcal{N}(0, \tau_j^2)$.

The two parameter-update proposals, (P1) and (P2), are symmetric. They only change the value of θ_j and leave $A^\dagger = A$ so that $\det \mathbf{D}(A^\dagger) = \det \mathbf{D}(A)$. From (2.9), one sees that the MH ratio is simply

$$\min\left\{1, \frac{f_{\mathbf{U}}(\mathbf{H}(\boldsymbol{\theta}^\dagger_A, \boldsymbol{\theta}^\dagger_I, A; \boldsymbol{\beta}))}{f_{\mathbf{U}}(\mathbf{H}(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A; \boldsymbol{\beta}))}\right\},$$

which can be computed very efficiently, especially for a normal error distribution. The proposal (P3) removes a variable from the active set and (P4) adds a variable to the active set. Both

propose moves between two subspaces. The two proposals are the reverse of each other and have a simple one-dimensional density. To be concrete, the MH ratio for proposal (P3) is

$$\min\left\{1, \frac{\pi(\boldsymbol{\theta}^{\dagger}_{A^{\dagger}}, \boldsymbol{\theta}^{\dagger}_{I^{\dagger}}, A^{\dagger})}{\pi(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A)} \cdot \frac{\phi(b_j; 0, \tau_j^2)}{1/2}\right\},$$

and analogously for proposal (P4). One needs to calculate the ratio between two determinants for these MH ratios,

$$\frac{|\det \mathbf{D}(A^{\dagger})|}{|\det \mathbf{D}(A)|} = \frac{\det \mathbf{C}_{A^{\dagger}A^{\dagger}}}{\det \mathbf{C}_{AA}}(w_j\lambda)^{|A|-|A^{\dagger}|} \tag{3.3}$$

by (2.7). As the two sets A and $A^{\dagger}$ differ by only one element, the ratio on the right-hand side can be calculated efficiently. When $|A|$ is large, we use the sweep operator to dynamically update $\mathbf{C}_{AA}^{-1}$ (the inverse of $\mathbf{C}_{AA}$) and obtain the ratio. See Appendix for further details. In general, however, a model-update proposal is more time-consuming than a parameter-update proposal.

This computational efficiency consideration motivates the following scheme in the MLS which uses both types of proposals. Let K be an integer between 1 and p and $\boldsymbol{\alpha} = (\alpha_j)_{1:p}$ be a vector with every $\alpha_j > 0$.

Routine 2 (MLS). Suppose the current state is $(\boldsymbol{\theta}^{(t)}, A^{(t)})$.

1. Draw K elements without replacement from $\{1, \ldots, p\}$ with the probability of drawing j proportional to α_j for each j. Let $M^{(t)}$ be the set of the K elements.
2. For $j \in M^{(t)}$, sequentially update each θ_j and the active set A by an MH step with a model-update proposal.
3. For $j \notin M^{(t)}$, sequentially update each θ_j by an MH step with a parameter-update proposal.

After the above p MH steps in an iteration, the state is updated to $(\boldsymbol{\theta}^{(t+1)}, A^{(t+1)})$.

The MLS has three input parameters, K, $\boldsymbol{\alpha}$, and $(\tau_j^2)_{1:p}$. Specification of these parameters that gives good empirical performance will be provided in the numerical examples (Section 3.6).

3.3 The Gibbs Lasso Sampler

Let $a_j = \mathbf{1}(j \in A)$ and $\mathbf{a} = (a_j)_{1:p}$. Conditional densities $\pi(\theta_j, a_j \mid \boldsymbol{\theta}_{-j}, \mathbf{a}_{-j})$ can be derived from the joint density π, which allows for the development of a Gibbs sampler. However, as each conditional sampling step involves calculation of one-dimensional integrals and sampling from truncated distributions, the Gibbs sampler is more time-consuming and less efficient than the MLS for all examples on which we have tested these algorithms.

3.4 Conditioning on Active Set

Suppose that we have constructed a Lasso-type estimate $\hat{\boldsymbol{\beta}}^*$ from an observed dataset and the set of selected variables is A^*, which defines an estimated model. One may want to study the sampling distribution of the estimator given the estimated model, that is, $[\hat{\boldsymbol{\beta}}_{A^*} \mid \mathcal{A} = A^*]$. Confidence intervals of penalized estimators have been constructed by approximating this distribution via local expansion of the ℓ_1 norm (Fan and Li 2001; Zou 2006). Since local approximation may not be accurate for a finite sample, Monte Carlo sampling from this conditional distribution may provide more accurate results. However, the direct sampling approach is not applicable in practice, because $\mathcal{A} = A^*$ is often a rare event unless p is very small. On the contrary, it is very efficient to draw samples by an MH algorithm from the conditional distribution

$$\pi(\mathbf{b}_{A^*}, \mathbf{s}_{I^*} \mid A^*) \propto f_{\mathbf{U}}(\mathbf{H}(\mathbf{b}_{A^*}, \mathbf{s}_{I^*}, A^*; \boldsymbol{\beta})), \tag{3.4}$$

where $I^* = \{1, \ldots, p\} \setminus A^*$, according to (2.12). The distribution of interest, $[\hat{\boldsymbol{\beta}}_{A^*} \mid \mathcal{A} = A^*]$, is a marginal distribution of (3.4). Since evaluation of this density does not involve calculation of determinants, each MH step is very fast.

Routine 3 (MLS given active set). Given the current state $(\mathbf{b}_{A^*}^{(t)}, \mathbf{s}_{I^*}^{(t)})$, sequentially draw $b_j^{(t+1)}$ for each $j \in A^*$ by an MH step with proposal (P1) and $s_j^{(t+1)}$ for each $j \notin A^*$ with proposal (P2) in one iteration.

3.5 Reparameterization View

To ease notation, write $\boldsymbol{\theta} = (\boldsymbol{\theta}_A, \boldsymbol{\theta}_I)$ and $\mathbf{H}_A(\boldsymbol{\theta}) = \mathbf{H}(\boldsymbol{\theta}, A) = \mathbf{H}(\boldsymbol{\theta}_A, \boldsymbol{\theta}_I, A; \boldsymbol{\beta})$ for $\boldsymbol{\theta}$ defined in (3.1). Suppose that $(\boldsymbol{\theta}^{(t)}, A^{(t)})$ are simulated by the MLS (Routine 2) and let $\mathbf{u}^{(t)} = \mathbf{H}(\boldsymbol{\theta}^{(t)}, A^{(t)})$. Since $\mathbf{H}$ is a bijection, $\mathbf{u}^{(t)}$ is a Markov chain that leaves $f_{\mathbf{U}}$ invariant. Therefore, the MLS can be understood as an MH algorithm targeting at $f_{\mathbf{U}}$ with moves designed under local reparameterization, $\boldsymbol{\theta} = \mathbf{H}_A^{-1}(\mathbf{u})$ for $\mathbf{u} \in U_A = \mathbf{H}_A(\Omega_A; \boldsymbol{\beta})$ (2.8). The Jacobian of this reparameterization is $[\mathbf{D}(A)]^{-1}$. Under this view, the MH ratio for a proposal $\mathbf{u}^{\dagger} = \mathbf{H}(\boldsymbol{\theta}^{\dagger}, A^{\dagger})$ given the current $\mathbf{u}$ is

$$\min\left\{1, \frac{f_{\mathbf{U}}(\mathbf{u}^{\dagger})}{f_{\mathbf{U}}(\mathbf{u})} \frac{q((\boldsymbol{\theta}^{\dagger}, A^{\dagger}), (\boldsymbol{\theta}, A))|\det \mathbf{D}(A)|^{-1}}{q((\boldsymbol{\theta}, A), (\boldsymbol{\theta}^{\dagger}, A^{\dagger}))|\det \mathbf{D}(A^{\dagger})|^{-1}}\right\},$$

which of course coincides with (3.2). Moreover, when $A = A^{\dagger}$ such as in proposals (P1) and (P2), the Jacobian determinants cancel out as we are using the same reparameterization $\mathbf{H}_A^{-1}$ for both the proposal and the current state. Otherwise, the ratio of the Jacobian determinants accounts for the use of different reparameterizations. Clearly, if $\mathbf{u}^{(t)} \sim f_{\mathbf{U}}$ then $(\boldsymbol{\theta}^{(t)}, A^{(t)}) \sim \pi$.

This view provides an insight into the computational efficiency of the MLS. Under normal error assumption, $f_{\mathbf{U}}$ is the density of a multivariate normal distribution, for which a simple MH algorithm is computationally tractable and can be quite efficient. We make a comparison with the direct sampler at a conceptual level. The direct sampler draws $\mathbf{U}$ via a linear transformation of $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, which costs n draws from a univariate normal distribution followed by a multiplication with a size $p \times n$ matrix. After that, we find $(\boldsymbol{\theta}, A) = \mathbf{H}^{-1}(\mathbf{U})$ by numerical minimization due to the lack of a closed-form inverse of the mapping $\mathbf{H}$. The MLS draws p ($< n$) univariate proposals in one iteration, and does not need any numerical procedure to map $\mathbf{u}^{(t)}$ back to the space Ω of $(\boldsymbol{\theta}, A)$ since the moves are by design in that space already. The mapping $\mathbf{H}$ from $(\boldsymbol{\theta}^{(t)}, A^{(t)})$ to $\mathbf{u}^{(t)}$ is simple and can be calculated analytically. This is fundamentally different from direct sampling which replies on a numerical procedure to find the image of each draw of $\mathbf{U}$ under the mapping $\mathbf{H}^{-1}$. The relatively time-consuming step in the MLS is calculating the ratio (3.3) when a model-update proposal is used, which can be done by at most sweeping a $|A| \times |A|$ matrix on a single position (Appendix). Owing to sparsity, $|A|$ is usually much smaller than p, which greatly speeds up this step. Since the target distribution in the space of $\mathbf{U}$ has a nice

Table 1. Simulated datasets for MCMC

Dataset	A	B	C	D
(n, p, σ^2)	(500, 100, 1)	(500, 200, 1)	(300, 100, 4)	(300, 200, 4)
$\lvert A^*\rvert$	23	22	25	57

unimodal density, the chain $\mathbf{u}^{(t)}$ often converges fast and has low autocorrelation. Consequently, we expect to see efficiency gain over direct sampling for the same amount of computing time, which will be confirmed numerically in the next subsection.

As in the following routine, by a special initialization such that $(\boldsymbol{\theta}^{(1)}, A^{(1)}) \sim \pi$, the MLS can reach equilibrium in one step, which totally removes the need for burn-in iterations. This will make our method suitable for parallel computing. See Section 7.4 for a more detailed discussion. However, to demonstrate the efficiency of the MLS as an independent method, we did not use Routine 4 in the numerical results.

Routine 4. Draw $(\mathbf{b}^{(1)}, \mathbf{s}^{(1)})$ from the direct sampler and let $(\boldsymbol{\theta}^{(1)}, A^{(1)})$ be its equivalent representation. With $(\boldsymbol{\theta}^{(1)}, A^{(1)})$ as the initial state, generate $(\boldsymbol{\theta}^{(t)}, A^{(t)})$ for $t = 2, \ldots, N$ by an MCMC algorithm targeting at π.

3.6 Numerical Examples

We demonstrate with numerical examples the effectiveness of the above MCMC algorithms by comparing against the direct sampling approach. To this end, we simulated four datasets with different combinations of n, p, and σ^2 (Table 1). The vector of true coefficients $\boldsymbol{\beta}_0$ has 10 nonzero components, $\beta_{0j} = 1$ for $j = 1, \ldots, 5$ and $\beta_{0j} = -1$ for $j = 6, \ldots, 10$. Each row of $\mathbf{X}$ was generated independently from $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma}_\mathbf{X})$, where the diagonal and the off-diagonal elements of $\boldsymbol{\Sigma}_\mathbf{X}$ are 1 and 0.25, respectively. Given the design matrix $\mathbf{X}$, the response vector $\mathbf{y}$ was drawn from $\mathcal{N}_n(\mathbf{X}\boldsymbol{\beta}_0, \sigma^2\mathbf{I}_n)$.

The weights w_j (1.2) were set to 1 for all the following numerical results. The Lars package by Hastie and Efron was applied to find the solution path for each dataset. The value of λ was chosen by minimizing the C_p criterion implemented in the package, which determined the estimated coefficients, $\hat{\boldsymbol{\beta}}^* = (\hat{\beta}_1^*, \ldots, \hat{\beta}_p^*)$, of a dataset. The number of selected variables, $|A^*|$, for each dataset is given in Table 1. We considered two types of error distributions, the normal distribution and the elliptically symmetric distribution. Correspondingly, we calculated $\hat{\sigma}^2$ by (2.20) with $\check{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^*$ or constructed $\hat{f}_\mathbf{U}$ by the approach in Section 2.5. For all the results, Step 2 of the direct sampler (Routine 1) was implemented with the Lars package.

We first examined the performance of the MLS on sampling from the joint distribution (2.9) given $\hat{\boldsymbol{\beta}}^*$ and $\hat{\sigma}^2$ or $\hat{f}_\mathbf{U}$. Let $\omega_j = \Phi(-|\hat{\beta}_j^*|/\zeta_j)$ for $j = 1, \ldots, p$, where ζ_j is the standard error of $\hat{\beta}_j^{\mathrm{OLS}}$ and Φ is the cumulative distribution function of $\mathcal{N}(0, 1)$. We set $K = p/5$ and $\alpha_j \propto \omega_j + \omega_0$, where $\omega_0 = \sum_j \omega_j/(5p)$ serves as a baseline weight so that each variable has a reasonable chance to be selected for model-update proposals. See Routine 2 for notations. Under this setting, if the estimate $\hat{\beta}_j^*$ is close to zero relative to ζ_j, it will have a higher chance for model-update proposals. The τ_j used in the proposals (Definition 1) was set to $2\zeta_j$. The MLS was applied to each dataset 10 times independently. Each run consisted of $L = 5,500$ iterations with the first 500 as the burn-in period. In what follows, the sampler is abbreviated as MLSn and MLSe under the normal and the elliptically symmetric error distributions, respectively.

Figure 1(a) is the scatterplot of the samples of $\hat{\beta}_1$ and $\hat{\beta}_{50}$, and illustrates that the distributions of some $\hat{\beta}_j$ indeed have a point mass at zero. The histogram of the subgradient S_{50} is shown in Figure 1(b) with two point masses on ± 1 and otherwise continuous. Mixing of the MLS was fast, as demonstrated with two chains in Figure 1(c), where the initial values were chosen to be about 20 standard deviations away from each other. Figure 1(d) shows the fast decay of the autocorrelation among the samples of a $\hat{\beta}_j$, decreasing to below 0.05 in 10 to 15 iterations. The acceptance rate of the model-update proposals was generally between 0.2 and 0.4. For the parameter-update proposals, the acceptance rate was between 0.2 and 0.4 for (P1) and was higher than 0.6 for (P2), which is an independent proposal.

From the MCMC samples, we estimated the selection probability $P_{s,j} = P(\hat{\beta}_j \neq 0)$, the 2.5% and the 97.5% quantiles of $\hat{\beta}_j$, and the mean and the standard deviation of the conditional distribution $[\hat{\beta}_j \mid \hat{\beta}_j \neq 0]$ for each j. Since theoretical values are not available, we applied the direct sampling approach to simulate 5,000 independent samples for each dataset under the normal error distribution. These independent samples were used to estimate the above quantities as the ground truth. The MSEs across 10 independent runs of the MLS were calculated, and reported in Table 2 are the average MSEs over all j for estimating the above five quantities. One clearly sees that all the estimates were very accurate. The MSE of the MLSe was greater than, but on the same order as, that of the MLSn for most estimates, which is expected due to the loss of efficiency without assuming a normal error distribution.

We compared the efficiency of the MLS against the direct sampler (DSn) under the same amount of running time and under the same normal error distribution. The DSn generated around 500 samples in the same amount of time for 5,500 iterations of the MLSn. The ratio of the MSE of the DSn to that of the MLSn was calculated for each estimate (Table 2). For most estimates, the MLSn seems to be more efficient and may reduce the MSE by 10% to 60%. The improvement was more significant for datasets A and B where the sample size $n = 500$. For the other two datasets, the MLSn showed a higher MSE in estimating selection probabilities but was more accurate for all other estimates. Furthermore, if the error distribution is more complicated such that one cannot simulate samples independently from the distribution, the efficiency of the direct sampler may be even lower. These results clearly confirm the notion that the MLS can serve as an efficient alternative to the direct sampling method for simulating from the sampling distribution of a Lasso-type estimator.

Next, we implemented Routine 3 to sample from the conditional distribution of $\hat{\boldsymbol{\beta}}$ given the model selected according to the C_p criterion, that is, $[\hat{\boldsymbol{\beta}}_{A^*} \mid \mathcal{A} = A^*]$ with $|A^*|$ given in Table 1. The same parameter setting as that in the previous example was used to run the MLSn and the MLSe. We estimated the 2.5% and the 97.5% quantiles, the mean, and the standard deviation of $\hat{\beta}_j$ for $j \in A^*$. The model space is composed of 2^p models, and the probability of the model A^*, $P(\mathcal{A} = A^*)$, is practically zero for the datasets used here. Therefore, the direct sampling approach is not applicable. This shows the advantage

Table 2. MSE comparison for simulation from the joint sampling distribution

Method	P_s	2.5%	97.5%	mean	SD
A					
MLSn	3.38×10^{-4}	1.82×10^{-5}	1.79×10^{-5}	4.36×10^{-6}	2.78×10^{-6}
MLSe	1.29	1.20	1.19	0.97	1.38
DSn	1.11	2.28	2.45	2.23	2.53
B					
MLSn	2.13×10^{-4}	2.89×10^{-5}	1.74×10^{-5}	1.22×10^{-5}	8.44×10^{-6}
MLSe	1.20	1.07	1.10	1.08	1.30
DSn	1.26	1.97	1.89	2.29	2.74
C					
MLSn	4.14×10^{-4}	1.23×10^{-4}	1.24×10^{-4}	3.20×10^{-5}	2.28×10^{-5}
MLSe	1.55	2.09	1.78	1.03	2.46
DSn	0.47	1.18	1.33	1.24	1.39
D					
MLSn	4.34×10^{-4}	2.96×10^{-4}	2.85×10^{-4}	6.37×10^{-5}	5.02×10^{-5}
MLSe	2.74	3.81	3.61	1.17	5.70
DSn	0.69	1.52	1.34	1.21	1.57

NOTE: For the MLSe and the DSn, reported is the ratio of MSE to that of the MLSn. The sweep operator was used in the MLS to calculate determinant ratios for dataset D.

and flexibility of the Monte Carlo algorithms. Since we cannot construct ground truth for this example, the accuracy of an estimate is measured by its variance across 10 independent runs of the MLS, averaging over $j \in A^*$ (Table 3). The variance of every estimate was on the order of 10^{-5} or smaller for datasets A, B, and C and was on the order of 10^{-4} or smaller for dataset D under both error models. This highlights the stability of the MLS in approximating sampling distributions across different runs. There were cases in which the variance of the MLSe was smaller. This does not necessarily suggest that the MLSe provided a more accurate estimate, as the loss of efficiency without the normal error assumption is likely to result in a higher bias.

4. HIGH-DIMENSIONAL SETTING

Recent efforts have established theoretical properties of ℓ_1-penalized linear regression in high dimension with $p > n$ (Meinshausen and Bühlmann 2006; Zhao and Yu 2006; Zhang and Huang 2008; Bickel, Ritov, and Tsybakov 2009). Under this setting, we assume $\operatorname{rank}(\mathbf{X}) = n < p$. Consequently, $\operatorname{row}(\mathbf{X})$ is an n-dimensional subspace of $\mathbb{R}^p$ and the Gram matrix $\mathbf{C}$ has n positive eigenvalues, denoted by $\Lambda_j > 0$, $j = 1, \ldots, n$. The associated orthonormal eigenvectors $\mathbf{v}_j \in \mathbb{R}^p$, $j = 1, \ldots, n$, form a basis for $\operatorname{row}(\mathbf{X})$. Choose orthonormal vectors $\mathbf{v}_{p+1}, \ldots, \mathbf{v}_p$ to form a basis for the null space of $\mathbf{X}$, $\operatorname{null}(\mathbf{X})$, and let $\mathbf{V} = (\mathbf{v}_1 | \ldots | \mathbf{v}_p)$. Then, $R = \{1, \ldots, n\}$ and $N = \{n+1, \ldots, p\}$ index the columns of $\mathbf{V}$ that form respective bases for $\operatorname{row}(\mathbf{X})$ and $\operatorname{null}(\mathbf{X})$.

Table 3. Variance comparison for conditional sampling given active set

Method	2.5%	97.5%	mean	SD
A				
MLSn	1.21×10^{-5}	1.28×10^{-5}	2.21×10^{-6}	1.03×10^{-6}
MLSe	0.90	1.02	0.92	1.05
B				
MLSn	1.47×10^{-5}	1.19×10^{-5}	3.19×10^{-6}	9.60×10^{-7}
MLSe	1.22	1.15	1.02	1.23
C				
MLSn	7.66×10^{-5}	8.65×10^{-5}	1.59×10^{-5}	7.08×10^{-6}
MLSe	1.07	0.95	1.01	1.00
D				
MLSn	1.67×10^{-4}	1.78×10^{-4}	2.55×10^{-5}	1.28×10^{-5}
MLSe	0.77	0.97	0.77	0.79

NOTE: Variance of the MLSe is reported as the ratio to that of the MLSn.

Assumption 1. Every n columns of $\mathbf{X}$ are linearly independent and every $(p - n)$ rows of $\mathbf{V}_N$ are linearly independent.

The first part of this assumption is sufficient for the columns of $\mathbf{X}$ being in general position, which guarantees that $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ is unique for any $\mathbf{y}$ and $\lambda > 0$ (Lemma 1). The second part will ease our derivation of the joint density of the augmented estimator. Note that Assumption 1 holds with probability one if the entries of $\mathbf{X}$ are drawn from a continuous distribution on $\mathbb{R}^{n \times p}$.

4.1 The Bijection

Although $\mathbf{U} = \frac{1}{n}\mathbf{X}^\mathsf{T}\varepsilon$ is a p-vector, by definition it always lies in $\operatorname{row}(\mathbf{X})$. Therefore, $\mathbf{V}_N^\mathsf{T}\mathbf{U} = \mathbf{0}$ and the n-vector $\mathbf{R} = \mathbf{V}_R^\mathsf{T}\mathbf{U}$ gives the coordinates of $\mathbf{U}$ with respect to the basis $\mathbf{V}_R$. If ε follows a continuous distribution on $\mathbb{R}^n$, then $\mathbf{R}$ has a proper density with respect to ξ_n. For example, if $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, then $\mathbf{R} \sim \mathcal{N}_n(\mathbf{0}, \frac{\sigma^2}{n}\boldsymbol{\Lambda})$ with $\boldsymbol{\Lambda} = \operatorname{diag}(\Lambda_1, \ldots, \Lambda_n)$. Now $\mathbf{R}$ plays the same role as $\mathbf{U}$ does in the low-dimensional case. We will use the known distribution of $\mathbf{R}$ to derive the distribution of the augmented estimator $(\hat{\boldsymbol{\beta}}_\mathcal{A}, \mathbf{S}_\mathcal{I}, \mathcal{A})$.

However, a technical difficulty is that when $p > n$, the map $\mathbf{H}$ defined in (2.6) is not a mapping from Ω to $\operatorname{row}(\mathbf{X})$ as $\mathbf{H}(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}) \in \mathbb{R}^p$ is not necessarily in $\operatorname{row}(\mathbf{X})$ for every $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega$. We thus need to remove those "illegal" points in Ω so that the image of $\mathbf{H}$ always lies in the row space of $\mathbf{X}$. This is achieved by imposing the constraint that

$$\mathbf{V}_N^\mathsf{T}\mathbf{H}(\hat{\boldsymbol{\beta}}_\mathcal{A}, \mathbf{S}_\mathcal{I}, \mathcal{A}; \boldsymbol{\beta}) = \mathbf{V}_N^\mathsf{T}\mathbf{U} = \mathbf{0}, \tag{4.1}$$

that is, the image of $\mathbf{H}$ must be orthogonal to null($\mathbf{X}$). It is more convenient to use the equivalent definition of $\mathbf{H}$ in (2.3), that is,

$$\mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}) = \mathbf{C}\hat{\boldsymbol{\beta}} + \lambda \mathbf{W}\mathbf{S} - \mathbf{C}\boldsymbol{\beta}. \tag{4.2}$$

Because $\mathbf{C}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) = \frac{1}{n}\mathbf{X}^{\mathsf{T}}\mathbf{X}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \in \text{row}(\mathbf{X})$, constraint (4.1) is equivalent to

$$\mathbf{V}_N^{\mathsf{T}}\mathbf{W}\mathbf{S} = \mathbf{V}_{\mathcal{A}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{A}\mathcal{A}}\mathbf{S}_{\mathcal{A}} + \mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}\mathbf{S}_{\mathcal{I}} = \mathbf{0}. \tag{4.3}$$

In words, the constraint is that the vector $\mathbf{W}\mathbf{S}$ must lie in row($\mathbf{X}$). Therefore, we have a more restricted space for the augmented estimator $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ in the high-dimensional case,

$$\Omega_r = \left\{(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega : \mathbf{V}_{AN}^{\mathsf{T}}\mathbf{W}_{AA}\text{sgn}(\mathbf{b}_A) + \mathbf{V}_{IN}^{\mathsf{T}}\mathbf{W}_{II}\mathbf{s}_I = \mathbf{0}\right\}. \tag{4.4}$$

Restricted to this space, $\mathbf{H}$ is a bijection.

Lemma 3. If $p > n$ and Assumption 1 holds, then for any $\boldsymbol{\beta}$ and $\lambda > 0$, the restriction of the mapping $\mathbf{H}$ (2.6) to Ω_r, denoted by $\mathbf{H}\,|_{\Omega_r}$, is a bijection that maps Ω_r onto row($\mathbf{X}$).

Proof. For any $\mathbf{U} \in \text{row}(\mathbf{X})$, there is a unique $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ such that $\mathbf{U} = \mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}) \in \text{row}(\mathbf{X})$ by Lemma 1. Thus, $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ satisfies the constraint (4.1) and lies in Ω_r. For any $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}) \in \Omega_r$, $\mathbf{V}_N^{\mathsf{T}}\mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}) = \mathbf{0}$ and $\mathbf{H}$ maps it into row($\mathbf{X}$). □

Remark 5. Fixing $\mathcal{A} = A$, (4.3) specifies $|N| = p - n$ constraints, and thus, the continuous components $(\hat{\boldsymbol{\beta}}_A, \mathbf{S}_I) \in \mathbb{R}^p$ lie in an n-dimensional subspace of Ω_A (2.8). The bijection $\mathbf{H}\,|_{\Omega_r}$ maps a finite number of n-dimensional subspaces onto row($\mathbf{X}$) which is an $\mathbb{R}^n$.

Now we represent the bijection $\mathbf{H}\,|_{\Omega_r}$ in terms of its coordinates with respect to $\mathbf{V}_R$ and equate it with $\mathbf{R} = \mathbf{V}_R^{\mathsf{T}}\mathbf{U}$:

$$\mathbf{R} = \mathbf{V}_R^{\mathsf{T}}\mathbf{H}(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}) \triangleq \mathbf{H}_r(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}; \boldsymbol{\beta}). \tag{4.5}$$

4.2 Joint Sampling Distribution

The distribution for $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A}) \in \Omega_r$ is completely given by the distribution of $\mathbf{R}$ via the bijective map $\mathbf{H}_r : \Omega_r \to \mathbb{R}^n$. The only task left is to determine the Jacobian of $\mathbf{H}_r$, taking into account the constraint (4.3). Left multiplying by $\mathbf{V}_R^{\mathsf{T}}$ both sides of Equation (2.5), with the simple facts that $\mathbf{V}_R^{\mathsf{T}}\mathbf{W}_{\mathcal{A}} = \mathbf{V}_{\mathcal{A}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{A}\mathcal{A}}$ and $\mathbf{V}_R^{\mathsf{T}}\mathbf{W}_{\mathcal{I}} = \mathbf{V}_{\mathcal{I}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$, gives

$$\mathbf{R} = \mathbf{V}_R^{\mathsf{T}}\mathbf{C}_{\mathcal{A}}\hat{\boldsymbol{\beta}}_{\mathcal{A}} + \lambda\mathbf{V}_{\mathcal{A}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{A}\mathcal{A}}\mathbf{S}_{\mathcal{A}} + \lambda\mathbf{V}_{\mathcal{I}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}\mathbf{S}_{\mathcal{I}} - \mathbf{V}_R^{\mathsf{T}}\mathbf{C}\boldsymbol{\beta}. \tag{4.6}$$

For any fixed value of $\mathcal{A}$, differentiating $\mathbf{R}$ and both sides of the constraint (4.3) with respect to $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}})$ give, respectively,

$$d\mathbf{R} = \mathbf{V}_R^{\mathsf{T}}\mathbf{C}_{\mathcal{A}}d\hat{\boldsymbol{\beta}}_{\mathcal{A}} + \lambda\mathbf{V}_{\mathcal{I}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}d\mathbf{S}_{\mathcal{I}}, \tag{4.7}$$

$$\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}d\mathbf{S}_{\mathcal{I}} = \mathbf{0}. \tag{4.8}$$

Therefore, the constraint implies that $d\mathbf{S}_{\mathcal{I}}$ is in null($\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$).

Lemma 4. If $p > n$ and Assumption 1 is satisfied, then the dimension of null($\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$) is $n - |\mathcal{A}| \geq 0$.

Proof. Under the assumption, the minimizer $\hat{\boldsymbol{\beta}}$ of (1.2) is unique and always has an active set with size $|\mathcal{A}| \leq \min\{n, p\} = n$. See Lemma 14 in Tibshirani (2013). If Assumption 1 is satisfied, any $|N|$ rows of $\mathbf{V}_N$ are linearly independent. Since $|\mathcal{I}| = p - |\mathcal{A}| \geq p - n = |N|$, the rank of the $|N| \times |\mathcal{I}|$ matrix, $\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$, is $p - n$. Then, it follows that the dimension of null($\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$) is $|\mathcal{I}| - (p - n) = n - |\mathcal{A}| \geq 0$. □

Let $\mathbf{B}(\mathcal{I}) \in \mathbb{R}^{|\mathcal{I}| \times (n - |\mathcal{A}|)}$ be an orthonormal basis for null($\mathbf{V}_{\mathcal{I}N}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}$). Let $d\tilde{\mathbf{S}}$ be the coordinates of $d\mathbf{S}_{\mathcal{I}}$ with respect to the basis $\mathbf{B}(\mathcal{I})$, that is, $d\mathbf{S}_{\mathcal{I}} = \mathbf{B}(\mathcal{I})d\tilde{\mathbf{S}}$. Note that $d\tilde{\mathbf{S}} \in \mathbb{R}^{n-|\mathcal{A}|}$ according to the above lemma. Then (4.7) becomes

$$\begin{aligned} d\mathbf{R} &= \mathbf{V}_R^{\mathsf{T}}\mathbf{C}_{\mathcal{A}}d\hat{\boldsymbol{\beta}}_{\mathcal{A}} + \lambda\mathbf{V}_{\mathcal{I}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}\mathbf{B}(\mathcal{I})d\tilde{\mathbf{S}} \\ &= \mathbf{T}(\mathcal{A})\begin{pmatrix} d\hat{\boldsymbol{\beta}}_{\mathcal{A}} \\ d\tilde{\mathbf{S}} \end{pmatrix}, \end{aligned} \tag{4.9}$$

where $\mathbf{T}(\mathcal{A}) = (\mathbf{V}_R^{\mathsf{T}}\mathbf{C}_{\mathcal{A}} \mid \lambda\mathbf{V}_{\mathcal{I}R}^{\mathsf{T}}\mathbf{W}_{\mathcal{I}\mathcal{I}}\mathbf{B}(\mathcal{I}))$, an $n \times n$ matrix, is the Jacobian of the map $\mathbf{H}_r$. The dimension of $(d\hat{\boldsymbol{\beta}}_{\mathcal{A}}, d\tilde{\mathbf{S}})$ is always n for any $\mathcal{A}$. This confirms the notion in Remark 5 that the continuous components $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}})$ lie in an n-dimensional subspace when $\mathcal{A}$ is fixed.

Now we are ready to derive the density for $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ in high dimension. For $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega_r$, $d\mathbf{s}_I = \mathbf{B}(I)d\tilde{\mathbf{s}}$ for some $d\tilde{\mathbf{s}} \in \mathbb{R}^{n-|A|}$ and $\xi_{n-|A|}(d\tilde{\mathbf{s}})$ gives the infinitesimal volume at $\mathbf{s}_I$ subject to constraint (4.4).

Theorem 2. Assume that $p > n$ and Assumption 1 holds. Let $f_{\mathbf{R}}$ be the probability density of $\mathbf{R}$ with respect to ξ_n. For $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega_r$, the joint distribution of $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$ is given by

$$\begin{aligned} &P(\hat{\boldsymbol{\beta}}_A \in d\mathbf{b}_A, \mathbf{S}_I \in d\mathbf{s}_I, \mathcal{A} = A) \\ &\quad = f_{\mathbf{R}}(\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}))|\det\mathbf{T}(A)|\xi_n(d\mathbf{b}_A d\tilde{\mathbf{s}}) \\ &\quad \triangleq \pi_r(\mathbf{b}_A, \mathbf{s}_I, A)\xi_n(d\mathbf{b}_A d\tilde{\mathbf{s}}). \end{aligned} \tag{4.10}$$

Particularly, if $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, then

$$\pi_r(\mathbf{b}_A, \mathbf{s}_I, A) = \phi_n\left(\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}); \mathbf{0}, n^{-1}\sigma^2\boldsymbol{\Lambda}\right)|\det\mathbf{T}(A)|. \tag{4.11}$$

Proof. The proof is analogous to that of Theorem 1. Let $\mathbf{r} = \mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}) \in \mathbb{R}^n$. For any fixed A,

$$d\mathbf{r} = \mathbf{T}(A)\begin{pmatrix} d\mathbf{b}_A \\ d\tilde{\mathbf{s}} \end{pmatrix}$$

from (4.9) and thus $\xi_n(d\mathbf{r}) = |\det\mathbf{T}(A)|\xi_n(d\mathbf{b}_A d\tilde{\mathbf{s}})$. With the bijective nature of $\mathbf{H}_r$ and its restriction to any A, a change of variable gives

$$\begin{aligned} P(\hat{\boldsymbol{\beta}}_A \in d\mathbf{b}_A, \mathbf{S}_I \in d\mathbf{s}_I, \mathcal{A} = A) &= P(\mathbf{R} \in d\mathbf{r}) \\ &= f_{\mathbf{R}}(\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}))|\det\mathbf{T}(A)|\xi_n(d\mathbf{b}_A d\tilde{\mathbf{s}}). \end{aligned}$$

If $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$ then $\mathbf{R} \sim \mathcal{N}_n(\mathbf{0}, \frac{\sigma^2}{n}\boldsymbol{\Lambda})$, which leads to (4.11) immediately. □

Remark 6. The density $\pi_r(\mathbf{b}_A, \mathbf{s}_I, A)$ does not depend on which orthonormal basis we choose for null($\mathbf{V}_{IN}^{\mathsf{T}}\mathbf{W}_{II}$). If $\mathbf{B}'(I)$ is another orthonormal basis, then $\mathbf{B}'(I) = \mathbf{B}(I)\mathbf{O}$, where $\mathbf{O}$ is an $(n - |A|) \times (n - |A|)$ orthogonal matrix and $|\det\mathbf{O}| = 1$. Correspondingly,

$$\mathbf{T}'(A) = (\mathbf{V}_R^{\mathsf{T}}\mathbf{C}_A \mid \lambda\mathbf{V}_{IR}^{\mathsf{T}}\mathbf{W}_{II}\mathbf{B}'(I)) = \mathbf{T}(A)\text{diag}(\mathbf{I}_{|A|}, \mathbf{O}),$$

and thus $|\det\mathbf{T}'(A)| = |\det\mathbf{T}(A)|$.

Remark 7. One may unify Theorems 1 and 2 with the use of cumbersome notations, but the idea is simple. Note that $\mathbf{T}(A)$

and $\mathbf{D}(A)$ in (2.6) are connected by

$$\mathbf{T}(A) = \mathbf{V}_R^{\mathsf{T}}(\mathbf{C}_A \mid \lambda \mathbf{W}_I \mathbf{B}(I)) = \mathbf{V}_R^{\mathsf{T}} \mathbf{D}(A) \operatorname{diag}(\mathbf{I}_{|A|}, \mathbf{B}(I)).$$

If $\operatorname{rank}(\mathbf{X}) = p \le n$, the set N reduces to the empty set and $\mathbf{V}_R = \mathbf{V}$. Hence, the constraint (4.3) no long exists, the space Ω_r is the same as Ω, and $\operatorname{null}(\mathbf{V}_{IN}^{\mathsf{T}} \mathbf{W}_{II})$ is simply $\mathbb{R}^{|I|}$ for any $I = \{1, \ldots, p\} \setminus A$. Choosing $\mathbf{B}(I) = \mathbf{I}_{|I|}$ leads to $d\mathbf{s}_I = d\tilde{\mathbf{s}}$, which shows that the probability in (4.10) reduces to that in (2.9). In this case, $\mathbf{T}(A) = \mathbf{V}^{\mathsf{T}} \mathbf{D}(A)$, and thus a column of $\mathbf{T}(A)$ gives the coordinates of the corresponding column of $\mathbf{D}(A)$ with respect to the basis $\mathbf{V}$.

In principle, one can develop MCMC algorithms to sample from the joint distribution (4.10). Development of such an algorithm is a little tedious because of the constraints in the definition of Ω_r and the use of different bases $\mathbf{B}(I)$ in different subspaces. However, the explicit density given in Theorem 2 allows us to develop very efficient IS algorithms for approximating tail probabilities with respect to the sampling distribution of a Lasso-type estimator.

5. *P*-VALUE CALCULATION BY IS

To simplify description, we focus on the high-dimensional setting with normal errors so that the distribution of interest is π_r (4.11). For a fixed $\mathbf{X}$, the density π_r, the bijection $\mathbf{H}_r$ (4.5) and the matrix $\mathbf{T}$ (4.9) are written as $\pi_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}, \sigma^2, \lambda)$, $\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}, \lambda)$, and $\mathbf{T}(A; \lambda)$, respectively, to explicitly indicate their dependency on different parameters. Suppose we are given a Lasso-type estimate $\hat{\boldsymbol{\beta}}^*$ for an observed dataset with a tuning parameter λ^*. Under the null model $\mathcal{H} : \boldsymbol{\beta} = \boldsymbol{\beta}_0, \sigma^2 = \sigma_0^2$, we want to calculate the p-value of some test statistic $T(\hat{\boldsymbol{\beta}}) \in \mathbb{R}$ constructed from the Lasso-type estimator $\hat{\boldsymbol{\beta}}$ for $\lambda = \lambda^*$. Precisely, the desired p-value is

$$\begin{aligned} q^* &= P(|T(\hat{\boldsymbol{\beta}})| \ge T^*; \mathcal{H}, \lambda^*) \\ &= \int_{\Omega_r^*} \pi_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}_0, \sigma_0^2, \lambda^*) \xi_n(d\mathbf{b}_A d\tilde{\mathbf{s}}), \end{aligned} \tag{5.1}$$

where $T^* = |T(\hat{\boldsymbol{\beta}}^*)|$ and $\Omega_r^* = \{(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega_r : |T(\mathbf{b})| \ge T^*\}$. Even if we can directly sample from $\pi_r(\bullet; \boldsymbol{\beta}_0, \sigma_0^2, \lambda^*)$, estimating q^* will be extremely difficult when it is very small. With the closed-form density π_r, we can use IS to solve this challenging problem.

5.1 Importance Sampling

Our target distribution is $\pi_r(\bullet; \boldsymbol{\beta}_0, \sigma_0^2, \lambda^*)$ and we propose to use $\pi_r(\bullet; \boldsymbol{\beta}_0, (\sigma^2)^\dagger, \lambda^\dagger)$ as a trial distribution to estimate expectations with respect to the target distribution via IS. First, note that the trial and the target distributions have the same support as the constraint in (4.4) that defines the space Ω_r only depends on $\mathbf{X}$. Thus, a sample from the trial distribution $\pi_r(\bullet; \boldsymbol{\beta}_0, (\sigma^2)^\dagger, \lambda^\dagger)$ also satisfies the constraint for the target distribution. Second, one can easily simulate from the trail distribution by the direct sampler (Routine 1). Third, the importance weight for a sample $(\mathbf{b}_A, \mathbf{s}_I, A)$ from the trial distribution can be calculated efficiently. Let $(r_1(\boldsymbol{\beta}, \lambda), \ldots, r_n(\boldsymbol{\beta}, \lambda)) = \mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}, \lambda) \in \mathbb{R}^n$ and note that $\det \mathbf{T}(A; \lambda) = \lambda^{n-|A|} \det \mathbf{T}(A; 1)$. Using the fact that $\boldsymbol{\Lambda} = \operatorname{diag}(\Lambda_1, \ldots, \Lambda_n)$, the importance weight

$$\begin{aligned} &\frac{\phi_n\left(\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}_0, \lambda^*); \mathbf{0}, n^{-1}\sigma_0^2 \boldsymbol{\Lambda}\right) |\det \mathbf{T}(A; \lambda^*)|}{\phi_n\left(\mathbf{H}_r(\mathbf{b}_A, \mathbf{s}_I, A; \boldsymbol{\beta}_0, \lambda^\dagger); \mathbf{0}, n^{-1}(\sigma^2)^\dagger \boldsymbol{\Lambda}\right) |\det \mathbf{T}(A; \lambda^\dagger)|} \\ &\quad \propto \exp\left[\frac{n}{2(\sigma^2)^\dagger} \sum_{i=1}^n \frac{r_i^2(\boldsymbol{\beta}_0, \lambda^\dagger)}{\Lambda_i} - \frac{n}{2\sigma_0^2} \sum_{i=1}^n \frac{r_i^2(\boldsymbol{\beta}_0, \lambda^*)}{\Lambda_i}\right] \\ &\quad \times \left(\frac{\lambda^*}{\lambda^\dagger}\right)^{n-|A|} \triangleq w(\mathbf{b}_A, \mathbf{s}_I, A; \sigma_0^2, \lambda^*). \end{aligned} \tag{5.2}$$

Essentially, for each sample, we only need to compute the image of the map $\mathbf{H}_r$ and two sums of squares.

Routine 5. Draw $(\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}$, $t = 1, \ldots, L$, from the trial distribution $\pi_r(\bullet; \boldsymbol{\beta}_0, (\sigma^2)^\dagger, \lambda^\dagger)$ by Routine 1. Then the IS estimate for the p-value q^* is given by

$$\hat{q}^{(\mathrm{IS})} = \frac{\sum_{t=1}^L w((\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}; \sigma_0^2, \lambda^*) \mathbf{1}(|T(\mathbf{b}^{(t)})| \ge T^*)}{\sum_{t=1}^L w((\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}; \sigma_0^2, \lambda^*)}. \tag{5.3}$$

The key is to choose the parameters $(\sigma^2)^\dagger$ and $\lambda^\dagger$ in the trial distribution so that we have a substantial fraction of samples for which $|T(\mathbf{b}^{(t)})| \ge T^*$. Next, we discuss some guidance on tuning these parameters.

5.2 Tuning Trial Distributions

We illustrate our procedure for tuning the trial distribution assuming $\boldsymbol{\beta}_0 = \mathbf{0}$, that is, the null hypothesis is $\mathcal{H}_0 : \boldsymbol{\beta} = \mathbf{0}, \sigma^2 = \sigma_0^2$. In this case, the problem is difficult when $P(\hat{\boldsymbol{\beta}} = \mathbf{0})$ is close to one under the target distribution $\pi_r(\bullet; \mathbf{0}, \sigma_0^2, \lambda^*)$. In other words, λ^* is too big to obtain any nonzero estimate of the coefficients and consequently the p-value $P(|T(\hat{\boldsymbol{\beta}})| \ge T^*; \mathcal{H}_0, \lambda^*)$ becomes a tail probability. Thus, one may want to choose the trial distribution $\pi_r(\bullet; \mathbf{0}, (\sigma^2)^\dagger, \lambda^\dagger)$ under which there is a higher probability for nonzero $\hat{\boldsymbol{\beta}}$. In general, we achieve this by choosing $(\sigma^2)^\dagger = M^\dagger \sigma_0^2$ $(M^\dagger > 1)$ and then tuning $\lambda^\dagger$ accordingly. When we increase σ^2, the variance of $\mathbf{U}$ increases and thus $\mathbf{U}$ will have a wider spread in row$(\mathbf{X})$. This will increase the variance of the augmented estimator $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$. As illustrated in Figure 2, a larger variance in $\mathbf{U}$ will lead to a more uniform distribution over different subspaces $\{\Omega_A\}$.

The following simple procedure is used to determine $\lambda^\dagger$ given $(\sigma^2)^\dagger$, which works very well based on our empirical study.

Routine 6. Draw $\mathbf{y}^{(t)}$ from $\mathcal{N}_n(\mathbf{0}, (\sigma^2)^\dagger \mathbf{I}_n)$ and calculate $\lambda^{(t)} = n^{-1} \|\mathbf{W}^{-1} \mathbf{X}^{\mathsf{T}} \mathbf{y}^{(t)}\|_\infty$ for $t = 1, \ldots, L_{\mathrm{pilot}}$. Then, set $\lambda^\dagger$ to the first quartile of $\{\lambda^{(t)} : t = 1, \ldots, L_{\mathrm{pilot}}\}$.

Setting $\hat{\boldsymbol{\beta}} = \mathbf{0}$ in (2.1), we have

$$n^{-1} \|\mathbf{W}^{-1} \mathbf{X}^{\mathsf{T}} \mathbf{y}\|_\infty = \lambda \|\mathbf{S}\|_\infty \le \lambda,$$

which shows that the $\lambda^{(t)}$ calculated in Routine 6 is the minimum value of λ with which $\hat{\boldsymbol{\beta}} = \mathbf{0}$ for $\mathbf{y}^{(t)}$. Therefore, under the trial distribution $\pi_r(\bullet; \mathbf{0}, (\sigma^2)^\dagger, \lambda^\dagger)$, $P(\hat{\boldsymbol{\beta}} = \mathbf{0})$ is around 25% and there is a 75% of chance for $\hat{\boldsymbol{\beta}}$ to have some nonzero components. This often results in a good balance between the dominating region of the target distribution $(\hat{\boldsymbol{\beta}} = \mathbf{0})$ and the region of interest Ω_r^* for p-value calculation (5.1). For all numerical examples in this article, we choose $M^\dagger = 5$ and $L_{\mathrm{pilot}} = 100$.

5.3 Multiple Tests

Consider multiple linear models with the same set of predictors,

$$\mathbf{y}_k = \mathbf{X}\boldsymbol{\beta}_k + \varepsilon_k, \quad k = 1, \ldots, m, \tag{5.4}$$

where $\mathbf{y}_k \in \mathbb{R}^n$, $\boldsymbol{\beta}_k \in \mathbb{R}^p$, and $\varepsilon_k \sim \mathcal{N}_n(\mathbf{0}, \sigma_k^2 \mathbf{I}_n)$. After proper rescaling of $\mathbf{y}_k$ and $\boldsymbol{\beta}_k$, we may assume that all σ_k^2 are identical, that is, $\sigma_k^2 = \sigma^2$. Suppose we are interested in testing against m null hypotheses $\mathcal{H}_k : \boldsymbol{\beta}_k = \mathbf{0}$ and $\sigma^2 = \sigma_0^2$, given Lasso-type estimates $\hat{\boldsymbol{\beta}}_k^*$ with $\lambda = \lambda_k^*$ for $k = 1, \ldots, m$. There are m p-values to calculate,

$$q_k^* = P(|T(\hat{\boldsymbol{\beta}})| \geq T_k^*; \mathcal{H}_k, \lambda_k^*), \tag{5.5}$$

where $T_k^* = |T(\hat{\boldsymbol{\beta}}_k^*)|$ for $k = 1, \ldots, m$. This problem occurs in various genomics applications. To give an example, $\mathbf{y}_k$ may be the expression level of gene k and $\mathbf{X}$ the expression levels of p transcription factors across n individuals. The transcription factors may potentially regulate the expression of a gene through the linear model (5.4). Rejection of $\mathcal{H}_k$ indicates that gene k is regulated by at least one of the p transcription factors.

To estimate all q_k^*, we only need to draw $(\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}$, $t = 1, \ldots, L$, from one trial distribution $\pi_r(\bullet; \mathbf{0}, (\sigma^2)^\dagger, \lambda^\dagger)$, in which $(\sigma^2)^\dagger = M^\dagger \sigma_0^2$ and $\lambda^\dagger$ is obtained by applying the same tuning procedure (Routine 6) once. Then, we calculate the importance weights by (5.2) for all target distributions, $\{w((\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}; \sigma_0^2, \lambda_k^*)\}_{1 \leq t \leq L}$, $k = 1, \ldots, m$, and construct estimates for all q_k^* by (5.3).

Remark 8. Alternatively, one may apply the Lars algorithm in the direct sampler to draw from the sampling distribution of $\hat{\boldsymbol{\beta}}$ given $\boldsymbol{\beta} = \mathbf{0}$ and $\sigma^2 = \sigma_0^2$ for all λ_k^*, as the Lars algorithm provides the whole solution path. The computing time of both methods is dominated by drawing samples and thus is comparable. However, when q_k^* is small, the IS method will be orders of magnitude more efficient than direct sampling, and when q_k^* is not too small, the accuracy of the two methods is on the same order. We will see this in the numerical examples. In addition, we do not have to use the Lars algorithm to draw from the trial distribution since there is no need to compute the solution path for IS. One thus has the freedom to choose other algorithms, such as coordinate descent (Friedman et al. 2007; Wu and Lange 2008), which may be more efficient when both n and p are large.

5.4 Numerical Examples

We first simulated two datasets to demonstrate the effectiveness in p-value calculation by the IS method for individual tests. Each row of $\mathbf{X}$ was generated from $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma}_\mathbf{X})$, where the diagonal and the off-diagonal elements of $\boldsymbol{\Sigma}_\mathbf{X}$ are 1 and 0.05, respectively. Given the predictors $\mathbf{X}$, the response vector $\mathbf{y}$ was drawn from $\mathcal{N}_n(\mathbf{X}\boldsymbol{\beta}_0, \sigma_0^2 \mathbf{I}_n)$. We set all weights $w_j = 1$. Table 4 reports the values of n, p, σ_0^2, and $\boldsymbol{\beta}_0$ for the two datasets. We applied the Lars algorithm on the two datasets and chose λ^* as the first λ along the solution path such that the Lasso estimate $\hat{\boldsymbol{\beta}}^*$ gave the correct number of active coefficients (Table 4). It turned out that $\hat{\boldsymbol{\beta}}^*$ only included one true active coefficient for both datasets. Let $A^* = \text{supp}(\hat{\boldsymbol{\beta}}^*)$ be the active set of $\hat{\boldsymbol{\beta}}^*$. We designed the following test statistics, $T_1 = \|\hat{\boldsymbol{\beta}}\|_1$, $T_2 = \|\hat{\boldsymbol{\beta}}\|_\infty$, and $\tilde{T}_j = |\hat{\beta}_j|$ for $j \in A^*$, and aimed to calculate p-values under the null hypothesis $\mathcal{H}_0 : \boldsymbol{\beta} = \mathbf{0}$ and $\sigma^2 = \sigma_0^2$.

Table 4. Simulated datasets for individual tests

Dataset	n	p	σ_0^2	$\boldsymbol{\beta}_0$	λ^*	$\lambda^\dagger$
E	5	10	1/4	$(2, -2, 0, \ldots, 0)$	1.65	0.60
F	10	20	1/4	$(1, 1, -1, -1, 0, \ldots, 0)$	0.315	0.57

We chose $(\sigma^2)^\dagger = 5\sigma_0^2$ and used Routine 6 to choose $\lambda^\dagger$ for the trial distributions. The values of $\lambda^\dagger$ for the two datasets are given in Table 4. When $(\sigma^2)^\dagger$ is sufficiently large for a dataset, the $\lambda^\dagger$ tuned by Routine 6 can be greater then λ^* (dataset F). The IS method (Routine 5) was applied with $L = 5,000$ to estimate p-values for all the above tests. This estimation procedure was repeated 10 times independently to obtain the standard deviation of an estimated p-value. We quantify the efficiency of an estimated p-value, $\hat{q}$, by its coefficient of variation $\text{cv}(\hat{q}) = \text{SD}(\hat{q})/\mathbb{E}(\hat{q})$, where the standard deviation and the mean are calculated across multiple runs. Table 5 summarizes the results, where $A^* = \{2, 3\}$ for dataset E and $A^* = \{4, 9, 15, 17\}$ for dataset F. One sees that the IS estimates were very accurate: Even for a tail probability as small as 10^{-21}, the coefficient of variation was less than or around 2. To benchmark the performance, we approximated the coefficient of variation of the estimate $\hat{q}^{(\text{DS})}$ constructed by direct sampling from the target distribution, $\text{cv}(\hat{q}^{(\text{DS})}) = \sqrt{(1-\bar{q})/(L\bar{q})}$ with $\bar{q} = \mathbb{E}(\hat{q}^{(\text{IS})})$. As reported in the table, for estimating an extremely small p-value, $\text{cv}(\hat{q}^{(\text{DS})})$ can be orders of magnitude greater than that of an IS estimate, and for a moderate p-value (around 10^{-2}), the two methods showed comparable performance.

Next, we simulated $m = 50$ datasets to test our p-value calculation for the multiple testing problem. We used the design matrix $\mathbf{X}$ in dataset F and $\sigma_0^2 = 1/4$. The response vector $\mathbf{y}_k$ was drawn from $\mathcal{N}_n(\mathbf{X}\boldsymbol{\beta}_k, \sigma_0^2 \mathbf{I}_n)$, where the true coefficient vector $\boldsymbol{\beta}_k$ is given in Table 6 for $k = 1, \ldots, 50$. For 10 datasets, $\boldsymbol{\beta}_k = \mathbf{0}$ and the null hypothesis is true. For 20 datasets, there are two large coefficients, which represents the case that the true model is sparse. The other 20 datasets mimic the scenario in which the true model has many relatively small coefficients. We chose λ_k^* as the first λ that gave two active coefficients along the solution path and used $T(\hat{\boldsymbol{\beta}}) = \|\hat{\boldsymbol{\beta}}\|_1$ as the test statistic. Summaries of

Table 5. Estimation of p-values for datasets E and F

	$\mathbb{E}(\hat{q}^{(\text{IS})})$	$\text{SD}(\hat{q}^{(\text{IS})})$	$\text{cv}(\hat{q}^{(\text{IS})})$	$\text{cv}(\hat{q}^{(\text{DS})})$
E				
T_1	3.7×10^{-19}	8.7×10^{-19}	2.37	2.32×10^{7}
T_2	5.7×10^{-15}	6.2×10^{-19}	1.09	1.88×10^{5}
$\tilde{T}_2$	6.3×10^{-21}	1.1×10^{-20}	1.81	1.79×10^{8}
$\tilde{T}_3$	8.4×10^{-15}	9.8×10^{-15}	1.16	1.54×10^{5}
F				
T_1	1.5×10^{-6}	4.5×10^{-7}	0.30	11.5
T_2	1.2×10^{-3}	1.2×10^{-4}	0.11	0.41
$\tilde{T}_4$	5.7×10^{-5}	2.2×10^{-5}	0.38	1.87
$\tilde{T}_9$	1.1×10^{-2}	1.3×10^{-3}	0.12	0.14
$\tilde{T}_{15}$	2.5×10^{-2}	1.9×10^{-3}	0.08	0.09
$\tilde{T}_{17}$	4.8×10^{-5}	1.5×10^{-5}	0.31	2.04

Table 6. Simulated datasets for multiple tests

Dataset	$\boldsymbol{\beta}_k$	Range of λ_k^*	Range of T_k^*
1–10	$(0, \ldots, 0)$	(0.16, 0.34)	(0.08, 0.23)
11–30	$(2, -2, 0, \ldots, 0)$	(0.88, 1.31)	(0.27, 0.84)
31–50	$(1/4, \ldots, 1/4)$	(0.70, 1.13)	(0.04, 0.51)

λ_k^* and $T_k^* = \|\hat{\boldsymbol{\beta}}_k^*\|_1$ for the 50 datasets are provided in Table 6 as well, from which we see that these datasets cover a wide range of λ_k^* and T_k^*. We chose $(\sigma^2)^\dagger = 5\sigma_0^2$. As seen from Routine 6, for identical $\mathbf{X}$ and $(\sigma^2)^\dagger$, the tuning procedure is the same. Therefore, we simply set $\lambda^\dagger = 0.57$, the value we used for dataset F (Table 4).

We simulated $L = 5,000$ samples from the trial distribution and estimated the p-values for all the 50 datasets. This procedure was repeated 10 times independently. The average over 10 runs of the estimated p-value, $\mathbb{E}(\hat{q}_k^{(\mathrm{IS})})$, is shown in Figure 3(a) for $k = 1, \ldots, 50$. As expected, most of the p-values for the first 10 datasets were not significant, while those for the other 40 datasets ranged from 10^{-4} to 10^{-30}, which confirms that $T(\hat{\boldsymbol{\beta}}) = \|\hat{\boldsymbol{\beta}}\|_1$ is a reasonable test statistic. Again, we see that even for p-values on the order of 10^{-30}, the coefficient of variation of an IS estimate was at most around 3 (Figure 3(b)). This provides huge gain in accuracy compared to direct sampling. Figure 3(c) plots $\log_{10}[\mathrm{cv}(\hat{q}_k^{(\mathrm{DS})})/\mathrm{cv}(\hat{q}_k^{(\mathrm{IS})})]$ for the 50 datasets, where $\mathrm{cv}(\hat{q}_k^{(\mathrm{DS})})$ was approximated in the same way as in the previous example. It is comforting to see that while the IS estimates $\hat{q}_k^{(\mathrm{IS})}$ showed huge improvement over the DS estimates in estimating a tail probability, they were only slightly worse than the DS estimates for an insignificant p-value. For the first 10 datasets, the coefficient of variation of $\hat{q}_k^{(\mathrm{IS})}$ was at most 7.9 times that of $\hat{q}_k^{(\mathrm{DS})}$. For majority of the other 40 datasets, the ratio of $\mathrm{cv}(\hat{q}_k^{(\mathrm{DS})})$ over $\mathrm{cv}(\hat{q}_k^{(\mathrm{IS})})$ was between 100 and 10^{10}.

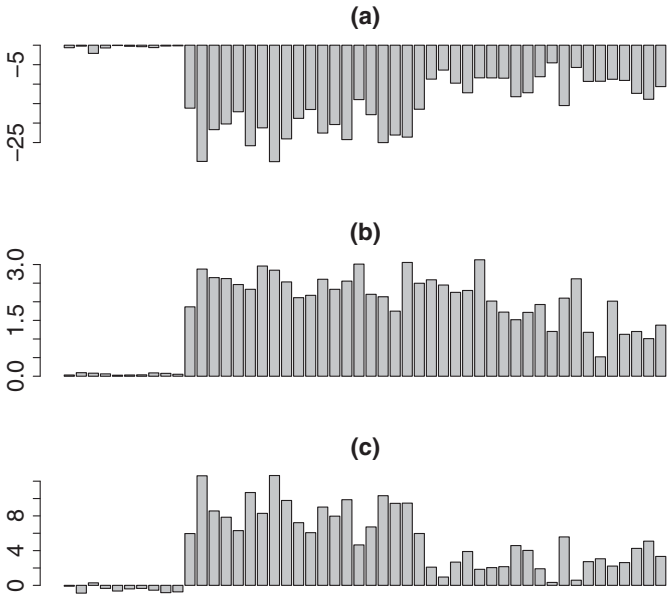

Figure 3. Estimation of p-values for 50 simulated datasets by IS with a single trial distribution: (a) $\log_{10}\mathbb{E}(\hat{q}_k^{(\mathrm{IS})})$, (b) $\mathrm{cv}(\hat{q}_k^{(\mathrm{IS})})$, and (c) $\log_{10}[\mathrm{cv}(\hat{q}_k^{(\mathrm{DS})})/\mathrm{cv}(\hat{q}_k^{(\mathrm{IS})})]$ for $k = 1, \ldots, 50$. Each bar in a plot gives the result for one dataset.

6. ESTIMATING SAMPLING DISTRIBUTIONS

For a vector $\mathbf{v} = (v_j)$, denote its active set by $\mathcal{A}(\mathbf{v}) = \{j : v_j \neq 0\}$. Let $\boldsymbol{\beta}_0 = (\beta_{0j})_{1:p}$ be the true coefficient vector and $A_0 = \mathcal{A}(\boldsymbol{\beta}_0)$. Consider model (1.1) with $\boldsymbol{\beta} = \boldsymbol{\beta}_0$ and $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$. The penalized loss in (1.2) is, up to an additive constant,

$$\begin{aligned}
&\frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 - \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_0\|_2^2 + n\lambda\sum_{j=1}^{p} w_j(|\beta_j| - |\beta_{0j}|) \\
&\quad = \frac{n}{2r_n^2}\boldsymbol{\delta}^\mathsf{T}\mathbf{C}\boldsymbol{\delta} - \frac{n}{r_n}\boldsymbol{\delta}^\mathsf{T}\mathbf{U} + n\lambda\sum_{j\in A_0} w_j\left(|\beta_{0j} + r_n^{-1}\delta_j| - |\beta_{0j}|\right) \\
&\qquad + \frac{n\lambda}{r_n}\sum_{j\notin A_0} w_j|\delta_j| \triangleq V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}), \qquad (6.1)
\end{aligned}$$

where $r_n > 0$, $\boldsymbol{\delta} = (\delta_j) = r_n(\boldsymbol{\beta} - \boldsymbol{\beta}_0)$, and $\mathbf{U} = \mathbf{X}^\mathsf{T}\varepsilon/n$. Assuming V has a unique minimizer,

$$\arg\min_{\boldsymbol{\delta}} V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}) = r_n(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \triangleq \hat{\boldsymbol{\delta}}, \qquad (6.2)$$

where $\hat{\boldsymbol{\beta}}$ is the Lasso-type estimator that minimizes (1.2). Suppose that $\check{\boldsymbol{\beta}}$ and $\hat{\sigma}$ are estimators of $\boldsymbol{\beta}_0$ and σ, following their respective sampling distributions. Let

$$\boldsymbol{\delta}^* = r_n(\boldsymbol{\beta}^* - \check{\boldsymbol{\beta}}) = \arg\min_{\boldsymbol{\delta}} V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{U}^*), \qquad (6.3)$$

where $\mathbf{U}^* = \mathbf{X}^\mathsf{T}\varepsilon^*/n$ and $\varepsilon^* \mid \hat{\sigma} \sim \mathcal{N}_n(\mathbf{0}, \hat{\sigma}^2\mathbf{I}_n)$. For random vectors $\mathbf{Z}_1$ and $\mathbf{Z}_2$, we use $\nu[\mathbf{Z}_1]$ to denote the distribution of $\mathbf{Z}_1$ and $\nu[\mathbf{Z}_1 \mid \mathbf{Z}_2]$ to denote the conditional distribution of $\mathbf{Z}_1$ given $\mathbf{Z}_2$. In general, $\nu[\mathbf{Z}_1 \mid \mathbf{Z}_2]$ is a random probability measure. The goal of this section is to derive nonasymptotic bounds on the difference between $\nu[\hat{\boldsymbol{\delta}}]$ and $\nu[\boldsymbol{\delta}^* \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$, with all proofs relegated to Section 8. Our results provide theoretical justifications for any method that estimates the uncertainty in $\hat{\boldsymbol{\beta}}$ via simulation from $\nu[\boldsymbol{\delta}^* \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$. In particular, for estimator augmentation, we draw $(\boldsymbol{\beta}^*, \mathbf{S}^*)|\check{\boldsymbol{\beta}}, \hat{\sigma}^2$ from an estimated sampling distribution whose density is given by (2.15) or (4.11) with $\boldsymbol{\beta} = \check{\boldsymbol{\beta}}$ and $\sigma^2 = \hat{\sigma}^2$.

6.1 Relevant Existing Results

We compile relevant published results here. The following restricted eigenvalue (RE) assumption is a special case of the one used in Lounici et al. (2011), which extends the original definition in Bickel, Ritov, and Tsybakov (2009).

Assumption 2 (RE(m, c_0)). For some positive integer $m \le p$ and a positive number c_0, the following condition holds:

$$\begin{aligned}
&\kappa(m, c_0) \\
&\quad \triangleq \min_{|A|\le m}\min_{\boldsymbol{\delta}\neq\mathbf{0}}\left\{\frac{\|\mathbf{X}\boldsymbol{\delta}\|_2}{\sqrt{n}\|\boldsymbol{\delta}_A\|_2} : \sum_{j\in A^c} w_j|\delta_j| \le c_0\sum_{j\in A} w_j|\delta_j|\right\} \\
&\quad > 0.
\end{aligned}$$

Lemma 5 is from Theorem 3.1 and its proof in Lounici et al. (2011), regarding the Lasso as a special case of the group Lasso (Yuan and Lin 2006) with the size of every group being one.

Lemma 5. Consider the model (1.1) with $\boldsymbol{\beta} = \boldsymbol{\beta}_0$ and $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, $\sigma^2 > 0$, and let $p \ge 2$, $n \ge 1$. Suppose that all the diagonal elements of $\mathbf{C}$ are 1 and $\mathbf{U} = (U_j)_{1:p} = \frac{1}{n}\mathbf{X}^\mathsf{T}\varepsilon$. Let $q_0 = |A_0| \le q$, where $1 \le q \le p$, and Assumption RE$(q, 3)$ be

satisfied. Choose $u > 1$ and

$$\lambda \geq \lambda_0 \triangleq \frac{2\sigma}{w_{\min}}\sqrt{(2+5u\log p)/n}, \tag{6.4}$$

where $w_{\min} = \inf\{w_j : j = 1, \ldots, p\}$. Then on the event $\mathcal{E} = \cap_{j=1}^{p}\{|U_j| \leq w_j\lambda/2\}$, which happens with probability at least $1 - 2p^{1-u}$, for any minimizer $\hat{\boldsymbol{\beta}}$ of (1.2) we have

$$|\mathcal{A}(\hat{\boldsymbol{\beta}})| \leq \frac{64\phi_{\max}}{\kappa^2(q,3)}\sum_{A_0}\frac{w_j^2}{w_{\min}^2}, \tag{6.5}$$

$$\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_1 \leq \frac{16\lambda}{\kappa^2(q,3)}\sum_{A_0}\frac{w_j^2}{w_{\min}}, \tag{6.6}$$

where $\phi_{\max}$ is the maximum eigenvalue of $\mathbf{C}$. If Assumption RE$(2q, 3)$ is satisfied, then on the same event,

$$\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_2 \leq \frac{4\sqrt{10}}{\kappa^2(2q,3)}\frac{\lambda\sum_{A_0} w_j^2}{w_{\min}\sqrt{q}} \triangleq \tau. \tag{6.7}$$

As an immediate consequence of this lemma, the Lasso-type estimator has the screening property assuming a suitable beta-min condition.

Lemma 6. Let the assumptions in Lemma 5 be satisfied. If $\inf_{A_0}|\beta_{0j}| > \tau$, then on the event $\mathcal{E}$, $A_0 \subseteq \mathcal{A}(\hat{\boldsymbol{\beta}})$ for any $\hat{\boldsymbol{\beta}}$. If $\inf_{A_0}|\beta_{0j}| > 2\tau$, then on the same event,

$$\{j : |\hat{\beta}_j| > \tau\} = A_0.$$

6.2 Known Variance

In this subsection, we assume that the noise variance σ^2 is known and fix $\hat{\sigma} = \sigma$ in the definition of $\boldsymbol{\delta}^*$ (6.3). We first regard $\check{\boldsymbol{\beta}} = (\check{\beta}_j)_{1:p}$ as a fixed vector and find conditions which are sufficient for the distribution of $\boldsymbol{\delta}^*$ to be close to that of $\hat{\boldsymbol{\delta}}$. Then, we construct an estimator that satisfies these conditions with high probability. Let

$$\eta \triangleq \sup_{j\in A_0}\frac{|\check{\beta}_j - \beta_{0j}|}{|\beta_{0j}|}. \tag{6.8}$$

Lemma 7. Assume that the columns of $\mathbf{X}$ are in general position. Fix $\hat{\sigma} = \sigma$. Suppose that $\check{\boldsymbol{\beta}}$ is a fixed vector in $\mathbb{R}^p$ so that (i) $\check{\beta}_j = 0$ for all $j \notin A_0$ and (ii) $\eta \in [0, 1)$ as defined in (6.8). Let $M_1 > 0$ and assume

$$\inf_{j\in A_0}|\beta_{0j}| > \frac{M_1}{r_n(1-\eta)}. \tag{6.9}$$

Then, we have

$$\nu[\boldsymbol{\delta}^* \mid \|\boldsymbol{\delta}^*\|_\infty < M_1] = \nu[\hat{\boldsymbol{\delta}} \mid \|\hat{\boldsymbol{\delta}}\|_\infty < M_1].$$

One possible way to construct $\check{\boldsymbol{\beta}}$ that satisfies with high probability conditions (i) and (ii) in Lemma 7 is to threshold the Lasso-type estimator $\hat{\boldsymbol{\beta}}$ by a constant $b_{\text{th}} > 0$, that is,

$$\check{\beta}_j = \hat{\beta}_j\mathbf{1}(|\hat{\beta}_j| > b_{\text{th}}), \quad j = 1, \ldots, p. \tag{6.10}$$

Theorem 3. Let the assumptions in Lemma 5 be satisfied and assume that the columns of $\mathbf{X}$ are in general position. Choose r_n such that $\|\hat{\boldsymbol{\delta}}\|_\infty < M_1$ with probability at least $1 - \alpha_1$. Fix $\hat{\sigma} = \sigma$, define $\check{\boldsymbol{\beta}}$ by (6.10) with $b_{\text{th}} = \tau$ (6.7), and assume

$$\inf_{j\in A_0}|\beta_{0j}| > \max\left\{2\tau, \frac{M_1}{r_n(1-\eta)}\right\}. \tag{6.11}$$

Then with probability at least $1 - 2p^{1-u}$, we have

$$\sup_{B\in\mathscr{R}^p}|P(\boldsymbol{\delta}^* \in B \mid \check{\boldsymbol{\beta}}) - P(\hat{\boldsymbol{\delta}} \in B)| \leq 2\alpha_1, \tag{6.12}$$

where $\mathscr{R}^p$ is the σ-field of p-dimensional Borel sets.

Remark 9. Depending on the estimator $\check{\boldsymbol{\beta}}$, the conditional probability $P(\boldsymbol{\delta}^* \in B \mid \check{\boldsymbol{\beta}})$ is a random variable. The probability $1 - 2p^{1-u}$ is with respect to the sampling distribution of $\mathbf{y}$. This theorem gives an explicit nonasymptotic bound (6.12) on the difference between the two probability measures, $\nu[\boldsymbol{\delta}^* \mid \check{\boldsymbol{\beta}}]$ and $\nu[\hat{\boldsymbol{\delta}}]$, and justifies simulation from the estimated sampling distribution with $\boldsymbol{\beta} = \check{\boldsymbol{\beta}}$.

Remark 10. Consider the asymptotic implications of Theorem 3 by allowing $q_0, p \to \infty$ as $n \to \infty$. Suppose that $w_{\min}$ and $w_{\max} = \sup_j w_j$ stay bounded away from 0 and ∞. For the Lasso, $w_{\min} = w_{\max} = 1$. Choose r_n so that $\|\hat{\boldsymbol{\delta}}\|_\infty \asymp_P 1$, that is, the convergence rate of $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_\infty$ is $1/r_n$. For any $\alpha_1 > 0$, there is $M_1 < \infty$ so that $P(\|\hat{\boldsymbol{\delta}}\|_\infty < M_1) \geq 1 - \alpha_1$. Let $q = q_0$ in Lemma 5 and assume that $\liminf_n \kappa(m, 3) > 0$ when $m = O(q_0) = o(n)$. With a suitable choice of $\lambda \asymp \sqrt{(\log p)/n}$,

$$\tau \asymp \sqrt{q_0}\lambda \asymp \sqrt{q_0(\log p)/n}.$$

Because $1 \asymp_P r_n\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_\infty = O_P(r_n\tau)$, we have $r_n^{-1} = O(\tau)$. Thus, the order of r_n satisfies

$$r_n^{-1} = O(\sqrt{q_0(\log p)/n}) \text{ and } r_n = O(\sqrt{n}). \tag{6.13}$$

Consequently, a sufficient condition for (6.11) is

$$\inf_{A_0}|\beta_{0j}| \gg \sqrt{q_0(\log p)/n} \tag{6.14}$$

with the order of b_{th} in between. Theorem 3 then implies that (6.12) holds with probability at least $1 - 2p^{1-u} \to 1$ as $p \to \infty$. Choosing α_1 arbitrarily close to zero, this demonstrates that

$$\sup_{B\in\mathscr{R}^p}|P(\boldsymbol{\delta}^* \in B \mid \check{\boldsymbol{\beta}}) - P(\hat{\boldsymbol{\delta}} \in B)| \xrightarrow{P} 0. \tag{6.15}$$

As $\inf_{A_0}|\beta_{0j}|$ may decay to zero at a rate slower than $\sqrt{q_0(\log p)/n} \to 0$, Theorem 3 applies in the high-dimensional setting ($p \gg n \to \infty$).

6.3 Unknown Variance

When σ^2 is unknown, recall that $\varepsilon^* \mid \hat{\sigma} \sim \mathcal{N}_n(\mathbf{0}, \hat{\sigma}^2\mathbf{I}_n)$ and $\mathbf{U}^* = \mathbf{X}^{\mathsf{T}}\varepsilon^*/n$. Denote the components of $\mathbf{U}^*$ by U_j^*, $j = 1, \ldots, p$. Define

$$\boldsymbol{\delta}_0^* = \arg\min_{\boldsymbol{\delta}} V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, (\sigma/\hat{\sigma})\mathbf{U}^*), \tag{6.16}$$

whose distribution does not depend on $\hat{\sigma}$ and is identical to that of $\boldsymbol{\delta}^*$ when $\hat{\sigma}$ is fixed to σ. Thus, results in Section 6.2 show that $\nu[\boldsymbol{\delta}_0^* \mid \check{\boldsymbol{\beta}}]$ is close to $\nu[\hat{\boldsymbol{\delta}}]$. We will further bound the difference, $\boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*$, in this subsection. In other words, $\boldsymbol{\delta}_0^*$ serves

as an intermediate variable between $\hat{\boldsymbol{\delta}}$ and $\boldsymbol{\delta}^*$ to help quantify the difference between their distributions.

Lemma 8. Assume that the columns of $\mathbf{X}$ are in general position. Let $\check{\boldsymbol{\beta}}$ and $\hat{\sigma}$ be fixed such that $|\mathcal{A}(\check{\boldsymbol{\beta}})| \le q$ with $1 \le q \le p$ and

$$\max\{|\hat{\sigma}/\sigma - 1|, |\sigma/\hat{\sigma} - 1|\} \le \zeta \in [0, 1). \tag{6.17}$$

If Assumption RE$(q, 3)$ is satisfied, then on the event $\cap_{j=1}^{p}\{|U_j^*| \le (1-\zeta)w_j\lambda/2\}$,

$$\frac{1}{n}\|\mathbf{X}(\boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*)\|_2^2 \le \frac{32 w_{\max} r_n^2 \lambda^2 \zeta}{w_{\min}\kappa^2(q,3)} \sum_{A_0} w_j^2. \tag{6.18}$$

Based on (6.18), we can obtain an upper bound on $\|\boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*\|_2$ via the restricted eigenvalues of $\mathbf{C}$. Define the minimum restricted eigenvalue of $\mathbf{C}$ for an integer $m \le p$ by

$$\phi_{\min}(m) = \min_{1 \le |\mathcal{A}(\mathbf{v})| \le m} \frac{\mathbf{v}^{\mathsf{T}}\mathbf{C}\mathbf{v}}{\|\mathbf{v}\|_2^2}, \tag{6.19}$$

and let

$$M_2 = \frac{128\phi_{\max}}{\kappa^2(q,3)} \sum_{A_0} \frac{w_j^2}{w_{\min}^2},$$

which is twice the upper bound on $|\mathcal{A}(\hat{\boldsymbol{\beta}})|$ in (6.5).

Theorem 4. Let the assumptions in Theorem 3 be satisfied but without fixing $\hat{\sigma}$ to σ. In addition, assume that (6.17) holds with probability at least $1 - \alpha_2$ and $\phi_{\min}(M_2) > 0$. Choose

$$\lambda \ge \frac{1+\zeta}{1-\zeta}\lambda_0, \tag{6.20}$$

where λ_0 is defined in (6.4). Then with probability at least $1 - (\alpha_2 + 2p^{1-u})$, we have

$$\sup_{B \in \mathscr{R}^p} |P(\boldsymbol{\delta}_0^* \in B \mid \check{\boldsymbol{\beta}}) - P(\hat{\boldsymbol{\delta}} \in B)| \le 2\alpha_1 \tag{6.21}$$

and

$$P\left\{\|\boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*\|_2^2 \le \frac{32 w_{\max} r_n^2 \lambda^2 \zeta \sum_{A_0} w_j^2}{w_{\min}\kappa^2(q,3)\phi_{\min}(M_2)} \,\middle|\, \check{\boldsymbol{\beta}}, \hat{\sigma}\right\} \ge 1 - 2p^{1-u}. \tag{6.22}$$

Remark 11. Assume that $\max\{|\hat{\sigma}/\sigma - 1|, |\sigma/\hat{\sigma} - 1|\} = O_P(\zeta_n)$ with $\zeta_n \to 0$. Following the asymptotic framework in Remark 10, we can establish by the same reasoning that

$$\sup_{B \in \mathscr{R}^p} |P(\boldsymbol{\delta}_0^* \in B \mid \check{\boldsymbol{\beta}}) - P(\hat{\boldsymbol{\delta}} \in B)| \xrightarrow{P} 0. \tag{6.23}$$

Suppose that $\liminf_n \phi_{\min}(m) > 0$ for $m = o(n)$ and that $\phi_{\max}$ is bounded from above by a constant or at least does not diverge too fast. Then, $M_2 = O(\phi_{\max} q_0) = o(n)$ and $\phi_{\min}(M_2)$ is bounded from below by a positive constant. The upper bound in (6.22) becomes $O(r_n^2\lambda^2\zeta_n q_0) = O(\zeta_n q_0 \log p)$ as $r_n = O(\sqrt{n})$ (6.13) and $\lambda^2 \asymp \log(p)/n$. If $\zeta_n = o(1/(q_0 \log p))$, then (6.22) implies that

$$P(\|\boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*\|_2 > \epsilon \mid \check{\boldsymbol{\beta}}, \hat{\sigma}) \to 0$$

for any $\epsilon > 0$. Combing with (6.23) this shows that, with probability tending to one, $\nu[\boldsymbol{\delta}_J^* \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$ converges weakly to $\nu[\hat{\boldsymbol{\delta}}_J]$ for any fixed index set $J \subseteq \{1, \ldots, p\}$. If σ is estimated by the scaled Lasso (Sun and Zhang 2012), one may reach $\zeta_n = O(q_0(\log p)/n)$ under certain conditions by their Theorem 2 and it is then sufficient to have $q_0 \log p \ll \sqrt{n}$. If p is fixed, we only need $\zeta_n = o(1)$ and hence any consistent estimator of σ will be sufficient. In this case, $r_n \asymp \sqrt{n}$ (6.13) and with probability tending to one, $\nu[\boldsymbol{\delta}^* \mid \check{\boldsymbol{\beta}}, \hat{\sigma}]$ converges weakly to $\nu[\hat{\boldsymbol{\delta}}]$.

The key assumptions on the underlying model are the RE assumption on the Gram matrix $\mathbf{C}$ and beta-min and sparsity assumptions on the true coefficients $\boldsymbol{\beta}_0$, which are comparable to those in Lemma 5 and Lemma 6. There is an extra assumption that $\phi_{\min}(M_2) > 0$ in Theorem 4, which is again imposed on the restricted eigenvalues of $\mathbf{C}$. If $\mathbf{X}$ is drawn from a continuous distribution on $\mathbb{R}^{n\times p}$, then with probability one $\phi_{\min}(m) > 0$ for any $m \le n$. It should be noted that we do not assume the irrepresentable condition (Meinshausen and Bühlmann 2006; Zhao and Yu 2006; Zou 2006), which is much stronger than the assumptions on the restricted eigenvalues of $\mathbf{C}$.

We compare our results to residual bootstrap for approximating the sampling distribution of the Lasso. To be precise, a residual bootstrap is equivalent to Routine 1 with $\boldsymbol{\beta}$ estimated by $\check{\boldsymbol{\beta}}$ and $\varepsilon^{(t)}$ drawn by resampling residuals. Assuming p is fixed, Knight and Fu (2000) argued that the residual bootstrap may be consistent if $\check{\boldsymbol{\beta}}$ is model selection consistent and Chatterjee and Lahiri (2011) establish such fixed-dimensional consistency when $\check{\boldsymbol{\beta}}$ is constructed by thresholding the Lasso, in the same spirit as (6.10). As discussed above, Theorem 4 applied to a fixed p is clearly in line with these previous works. However, our results are much more general by providing explicit nonasymptotic bounds that imply consistency when $p \gg n \to \infty$. More recently, Chatterjee and Lahiri (2013) have shown that the residual bootstrap is consistent for the adaptive Lasso (Zou 2006) when $p > n \to \infty$ under a number of conditions. A fundamental difference is that the weights w_j are specified by an initial $\sqrt{n}$-consistent estimator in their work and will not stay bounded as $n \to \infty$. Therefore, their results do not apply to the Lasso. On the contrary, the results in this section are derived assuming the weights w_j are constants without being specified by any initial estimator. In addition, Chatterjee and Lahiri (2013) imposed in their Theorem 5.1 that $\inf_{A_0} |\beta_{0j}| \ge K$ for some $K \in (0, \infty)$ when $p > n$, which disallows the decay of the magnitudes of nonzero coefficients as n grows. This is considerably stronger than our assumption (6.14).

7. GENERALIZATIONS AND DISCUSSIONS

7.1 Random Design

We generalize the Monte Carlo methods to a random design, assuming that $\mathbf{X}$ is drawn from a distribution $f_{\mathbf{X}}$. The distribution of the augmented estimator $(\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}, \mathcal{A})$, (2.9) and (4.10), becomes a conditional distribution given $\mathbf{X} = \mathbf{x}$, written as $\pi(\mathbf{b}_A, \mathbf{s}_I, A \mid \mathbf{x})$ and $\pi_r(\mathbf{b}_A, \mathbf{s}_I, A \mid \mathbf{x})$, respectively.

In the low-dimensional setting, we may generalize the MLS (Routine 2) to draw samples from $\pi(\mathbf{b}_A, \mathbf{s}_I, A, \mathbf{x}) = \pi(\mathbf{b}_A, \mathbf{s}_I, A \mid \mathbf{x}) f_{\mathbf{X}}(\mathbf{x})$ and approximate the sampling distribution of $\hat{\boldsymbol{\beta}}$. It may be difficult to assume or estimate a reliable density for $\mathbf{X}$, but it is sufficient for the development of an MH sampler under a random design (rdMLS) if we can draw from

$f_{\mathbf{X}}(\mathbf{x})$. As seen below, we do not need an explicit form of $f_{\mathbf{X}}(\mathbf{x})$ for computing the MH ratio (7.1) and thus may draw $\mathbf{x}^\dagger$ by the bootstrap.

Routine 7 (rdMLS). Suppose the current sample is $(\mathbf{b}_A, \mathbf{s}_I, A, \mathbf{x})^{(t)}$.

1. Draw $\mathbf{x}^\dagger$ from $f_{\mathbf{X}}$, and accept it as $\mathbf{x}^{(t+1)}$ with probability
$$\min\left\{1, \frac{\pi((\mathbf{b}_A, \mathbf{s}_I, A)^{(t)} \mid \mathbf{x}^\dagger)}{\pi((\mathbf{b}_A, \mathbf{s}_I, A)^{(t)} \mid \mathbf{x}^{(t)})}\right\}; \tag{7.1}$$
otherwise, set $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)}$.
2. Regarding $\pi(\mathbf{b}_A, \mathbf{s}_I, A \mid \mathbf{x}^{(t+1)})$ as the target density, apply one iteration of the MLS (Routine 2) to obtain $(\mathbf{b}_A, \mathbf{s}_I, A)^{(t+1)}$.

Generalization of the IS algorithm (Routine 5) is also straightforward. Draw $\mathbf{x}^{(t)}$ from $f_{\mathbf{X}}$ and draw $(\mathbf{b}_A, \mathbf{s}_I, A)^{(t)}$ from the trial distribution given $\mathbf{X} = \mathbf{x}^{(t)}$. Calculate importance weights by (5.2) with $\mathbf{X} = \mathbf{x}^{(t)}$, and apply the same estimation (5.3). Again, an explicit expression for $f_{\mathbf{X}}$ is unnecessary. But bootstrap sampling from $\mathbf{X}$ is not a choice for the high-dimensional setting, because a bootstrap sample from $\mathbf{X}$ violates Assumption 1.

7.2 Model Selection Consistency

The distribution of the augmented estimator may help establish asymptotic properties of a Lasso-type estimator. Here, we demonstrate this point by studying the model selection consistency of the Lasso. Our goal is not to establish new asymptotic results, but to provide an intuitive and geometric understanding of the technical conditions in existing work. Recall that $\mathcal{A}$ and A_0 are the respective active sets of $\hat{\boldsymbol{\beta}}$ and $\boldsymbol{\beta}_0$. Let $q_0 = |A_0|$ and $\mathbf{s}_0 = \operatorname{sgn}(\boldsymbol{\beta}_{0A_0})$. Without loss of generality, assume $A_0 = \{1, \ldots, q_0\}$ and $I_0 = \{q_0 + 1, \ldots, p\}$. We allow both p and q_0 to grow with n.

Definition 2 (sign consistency (Meinshausen and Yu 2009)). We say that $\hat{\boldsymbol{\beta}}$ is sign consistent for $\boldsymbol{\beta}_0$ if
$$P(\mathcal{A} = A_0, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_{A_0}) = \mathbf{s}_0) \to 1, \text{ as } n \to \infty. \tag{7.2}$$

If $\hat{\boldsymbol{\beta}}$ is unique, the size of its active set $|\mathcal{A}| \le n$ (Lemma 4), and thus $\mathbf{D}(\mathcal{A})$ is invertible from (2.7). Therefore, the definitions of $\boldsymbol{\mu}(A, \mathbf{s}_A; \boldsymbol{\beta})$ and $\boldsymbol{\Sigma}(A; \sigma^2)$ in (2.13) and (2.14) are also valid for any $(\mathbf{b}_A, \mathbf{s}_I, A) \in \Omega_r$ when $p > n$. Rewrite the KKT condition (2.6) as
$$\boldsymbol{\Theta} = [\mathbf{D}(\mathcal{A})]^{-1}\mathbf{U} + \boldsymbol{\mu}(\mathcal{A}, \mathbf{S}_{\mathcal{A}}; \boldsymbol{\beta}_0), \tag{7.3}$$
where $\boldsymbol{\Theta} = (\hat{\boldsymbol{\beta}}_{\mathcal{A}}, \mathbf{S}_{\mathcal{I}}) \in \mathbb{R}^p$. Fixing $\mathcal{A} = A_0$ and $\mathbf{S}_{A_0} = \mathbf{s}_0$ in (7.3), we define a random vector
$$\mathbf{Z} = [\mathbf{D}(A_0)]^{-1}\mathbf{U} + \boldsymbol{\mu}(A_0, \mathbf{s}_0; \boldsymbol{\beta}_0) \tag{7.4}$$
via an affine map of $\mathbf{U}$. Note that we always have $\mathbb{E}(\mathbf{U}) = \mathbf{0}$ and $\operatorname{var}(\mathbf{U}) = \frac{\sigma^2}{n}\mathbf{C} \ge 0$, regardless of the sizes of n and p. When $p > n$, $\operatorname{var}(\mathbf{U})$ is semipositive definite, meaning that components of $\mathbf{U}$ are linearly dependent of each other, since $\mathbf{U}$ only lies in $\operatorname{row}(\mathbf{X})$, a proper subspace of $\mathbb{R}^p$. Consequently, $\mathbb{E}(\mathbf{Z}) = \boldsymbol{\mu}(A_0, \mathbf{s}_0; \boldsymbol{\beta}_0) \triangleq \boldsymbol{\mu}^0$ and $\operatorname{var}(\mathbf{Z}) = \boldsymbol{\Sigma}(A_0; \sigma^2) \triangleq \boldsymbol{\Sigma}^0 \ge 0$. Simple calculation from (2.13) and (2.14) gives
$$\begin{pmatrix} \boldsymbol{\mu}^0_{A_0} \\ \boldsymbol{\mu}^0_{I_0} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\beta}_{0A_0} - \lambda \mathbf{C}^{-1}_{A_0A_0}\mathbf{W}_{A_0A_0}\mathbf{s}_0 \\ \mathbf{W}^{-1}_{I_0I_0}\mathbf{C}_{I_0A_0}\mathbf{C}^{-1}_{A_0A_0}\mathbf{W}_{A_0A_0}\mathbf{s}_0 \end{pmatrix}, \tag{7.5}$$
$$\boldsymbol{\Sigma}^0 = \frac{\sigma^2}{n}\begin{pmatrix} \mathbf{C}^{-1}_{A_0A_0} & \mathbf{0} \\ \mathbf{0} & \lambda^{-2}\mathbf{W}^{-1}_{I_0I_0}\mathbf{C}_{I_0|A_0}\mathbf{W}^{-1}_{I_0I_0} \end{pmatrix}, \tag{7.6}$$
where $\mathbf{C}_{I_0|A_0} = \mathbf{C}_{I_0I_0} - \mathbf{C}_{I_0A_0}\mathbf{C}^{-1}_{A_0A_0}\mathbf{C}_{A_0I_0}$.

Lemma 9. If the columns of $\mathbf{X}$ are in general position and ε is iid with mean zero and variance σ^2, then for any $p \ge 1$ and $n \ge 1$,
$$P(\mathcal{A} = A_0, \operatorname{sgn}(\hat{\boldsymbol{\beta}}_{A_0}) = \mathbf{s}_0) = P(\mathbf{Z} \in \Omega_{A_0,\mathbf{s}_0}), \tag{7.7}$$
where $\Omega_{A_0,\mathbf{s}_0}$ is defined by (2.17).

Proof. If $\mathcal{A} = A_0$ and $\mathbf{S}_{A_0} = \mathbf{s}_0$, then $\boldsymbol{\Theta} \in \Omega_{A_0,\mathbf{s}_0}$ by definition. In this case, (7.3) reduces to (7.4) and we have $\mathbf{Z} = \boldsymbol{\Theta} \in \Omega_{A_0,\mathbf{s}_0}$. Reversely, if $\mathbf{Z} \in \Omega_{A_0,\mathbf{s}_0}$, then $(\mathbf{Z}, A_0)$ is a solution to the KKT condition (7.3). By uniqueness, $(\boldsymbol{\Theta}, \mathcal{A}) = (\mathbf{Z}, A_0)$ and therefore $\mathcal{A} = A_0$ and $\mathbf{S}_{A_0} = \mathbf{s}_0$. □

Consequently, to establish sign consistency, we only need a set of sufficient conditions for $P(\mathbf{Z} \in \Omega_{A_0,\mathbf{s}_0}) \to 1$: (C1) $\operatorname{sgn}(\boldsymbol{\mu}^0_{A_0}) = \mathbf{s}_0$. (C2) $\|\boldsymbol{\mu}^0_{I_0}\|_\infty \le c$ for some $c \in (0, 1)$. (C3) Let Z_j and μ^0_j be the jth components of $\mathbf{Z}$ and $\boldsymbol{\mu}^0$, respectively. As $n \to \infty$,
$$P\left(|Z_j - \mu^0_j| < \delta_j, \ \forall j\right) \to 1, \tag{7.8}$$
where $\delta_j = |\mu^0_j|$ for $j \in A_0$ and $\delta_j = 1 - c$ for $j \in I_0$.

The first two conditions ensure that $\boldsymbol{\mu}^0 = \mathbb{E}(\mathbf{Z})$ lies in the interior of $\Omega_{A_0,\mathbf{s}_0}$. The third condition guarantees that $\mathbf{Z}$ always stays in a box centered at $\boldsymbol{\mu}^0$, and the box is contained in $\Omega_{A_0,\mathbf{s}_0}$ if (C1) and (C2) hold. These conditions have a simple and intuitive geometric interpretation illustrated in Figure 4.

Lemma 10. Assume the columns of $\mathbf{X}$ are in general position and ε is iid with mean zero and variance σ^2. If conditions (C1), (C2), and (C3) hold, then $\hat{\boldsymbol{\beta}}$ is sign consistent for $\boldsymbol{\beta}_0$, regardless of the relative size between p and n.

Now we may recover some of the conditions for establishing consistency of the Lasso in the literature. In what follows, let

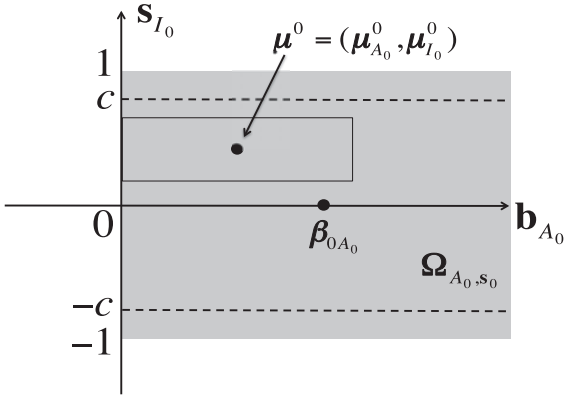

Figure 4. Geometric interpretation of the conditions for sign consistency. Shaded area represents $\Omega_{A_0,\mathbf{s}_0}$, where $\mathbf{s}_0 = \operatorname{sgn}(\boldsymbol{\beta}_{0A_0})$.

$\mathbf{W} = \mathbf{I}_p$ in (7.5) and (7.6). Condition (C2) is the strong irrepresentable condition (Meinshausen and Bühlmann 2006; Zhao and Yu 2006; Zou 2006):

$$\|\mathbf{C}_{I_0A_0}\mathbf{C}_{A_0A_0}^{-1}\mathbf{s}_0\|_\infty \le c \in (0,1).$$

Assume that the minimum eigenvalue of $\mathbf{C}_{A_0A_0}$ is bounded from below by $\phi_0 > 0$, which is equivalent to condition (6) in Zhao and Yu (2006). Condition (C1) holds if

$$\frac{\lambda\|\mathbf{C}_{A_0A_0}^{-1}\mathbf{s}_0\|_\infty}{\inf_{j\in A_0}|\beta_{0j}|} \le \frac{\lambda\phi_0^{-1}\|\mathbf{s}_0\|_2}{\inf_{j\in A_0}|\beta_{0j}|} = \frac{\phi_0^{-1}\lambda\sqrt{q_0}}{\inf_{j\in A_0}|\beta_{0j}|} \to 0. \quad (7.9)$$

This shows that some version of a beta-min condition is necessary to enforce a lower bound for $\inf_{j\in A_0}|\beta_{0j}|$. For example, we may assume that

$$\lim_{n\to\infty} n^{a_1} \inf_{j\in A_0}|\beta_{0j}| \ge M_3, \quad (7.10)$$

for some positive constants M_3 and a_1, which is the same as (8) in Zhao and Yu (2006). Then, one needs to choose $\lambda = o(n^{-a_1}/\sqrt{q_0}) \to 0$ for (C1) to hold.

Let $\mathbf{d} = (d_1, \ldots, d_p)$ such that $\mathbf{d}_{A_0} = \sigma^2\mathrm{diag}(\mathbf{C}_{A_0A_0}^{-1})$ and $\mathbf{d}_{I_0} = \sigma^2\mathrm{diag}(\mathbf{C}_{I_0|A_0})$. To establish condition (C3), assume that $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$ and $\mathrm{diag}(\mathbf{C})$ is bounded from above. Then, all d_j are bounded and let $d_* < \infty$ be an upper bound of $\{d_1, \ldots, d_p\}$. Furthermore, Z_j follows a univariate normal distribution: For $j \in A_0$, $Z_j \sim \mathcal{N}(\mu_j^0, n^{-1}d_j)$ and for $j \in I_0$, $Z_j \sim \mathcal{N}(\mu_j^0, n^{-1}\lambda^{-2}d_j)$ according to (7.6) with $\mathbf{W} = \mathbf{I}_p$. By (7.9) and (7.10), $\delta_j = |\mu_j^0| \ge \frac{1}{2}M_3n^{-a_1}$ for $j \in A_0$ as $n \to \infty$, and

$$P\left(\sup_{j\in A_0}|Z_j - \mu_j^0| \ge \frac{1}{2}M_3n^{-a_1}\right) \le 2q_0\exp\left(-\frac{M_3^2n^{1-2a_1}}{8d_*}\right) \to 0, \quad (7.11)$$

as long as $a_1 < 1/2$ and $q_0 < n$. Since $\delta_j = 1 - c$ for all $j \in I_0$,

$$P\left(\sup_{j\in I_0}|Z_j - \mu_j^0| \ge 1 - c\right) \le 2\exp\left(-\frac{n\lambda^2(1-c)^2}{2d_*} + \log p\right) \to 0, \quad (7.12)$$

if $(\log p)/(n\lambda^2) \to 0$ and $n\lambda^2 \to \infty$. Clearly, the above two inequalities imply (7.8). Therefore, λ must satisfy $\sqrt{(\log p)/n} \ll \lambda = o(n^{-a_1}/\sqrt{q_0})$, which implies that $\sqrt{q_0(\log p)/n} = o(n^{-a_1})$. This is consistent with the beta-min condition in Meinshausen and Bühlmann (2006): $\inf_{j\in A_0}|\beta_{0j}| \gg \sqrt{q_0(\log p)/n}$. In summary, choosing $a_1, a_2, a_3 > 0$ such that $a_2 + a_3 < 1 - 2a_1$, the Lasso can be consistent for model selection with $q_0 = O(n^{a_2})$ and $p = O(\exp(n^{a_3}))$, both diverging with n. For more general scaling of (n, p, q_0), see the work by Wainwright (2009).

Remark 12. The term $\log p$ in (7.12) can be replaced by $\log(p - q_0)$, which will improve the bound if q_0/p does not vanish as $n \to \infty$. Moreover, both inequalities (7.11) and (7.12) are applicable to sub-Gaussian noise.

7.3 Bayesian Interpretation

It is well-known that the Lasso can be interpreted as the mode of the posterior distribution of $\boldsymbol{\beta}$ under a Laplace prior. However, the posterior distribution itself is continuous on $\mathbb{R}^p$. If we draw $\boldsymbol{\beta}$ from this posterior distribution, every component of $\boldsymbol{\beta}$ will be nonzero with probability one. In this sense, sampling from this posterior distribution does not provide a direct solution to model selection, which seems unsatisfactory from a Bayesian perspective. Here, we discuss a different Bayesian interpretation of the Lasso-type estimator $\hat{\boldsymbol{\beta}}$ from a sampling distribution point of view.

Assume that $\mathrm{rank}(\mathbf{X}) = p < n$ and thus $\mathbf{C}$ is invertible. Under the noninformative prior $p(\boldsymbol{\beta}, \sigma^2) \propto 1/\sigma^2$ and the assumption that $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$, the conditional and marginal posterior distributions of $\boldsymbol{\beta}$ are

$$\boldsymbol{\beta} \mid \sigma^2, \mathbf{y} \sim \mathcal{N}_p\big(\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}, n^{-1}\sigma^2\mathbf{C}^{-1}\big), \quad (7.13)$$

$$\boldsymbol{\beta} \mid \mathbf{y} \sim t_{n-p}\big(\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}, n^{-1}\hat{\sigma}^2\mathbf{C}^{-1}\big), \quad (7.14)$$

where $\hat{\sigma}^2$ is given by (2.20) with $\check{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ and $t_{n-p}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is the multivariate t distribution with $(n - p)$ degrees of freedom, location $\boldsymbol{\mu}$, and scale matrix $\boldsymbol{\Sigma}$.

Following the decision theory framework, let $\boldsymbol{\eta} \in \mathbb{R}^p$ be a decision regarding $\boldsymbol{\beta}$ that incurs the loss

$$\ell_B(\boldsymbol{\eta}, \boldsymbol{\beta}) = \frac{1}{2}(\boldsymbol{\eta} - \boldsymbol{\beta})^{\mathsf{T}}\mathbf{C}(\boldsymbol{\eta} - \boldsymbol{\beta}) + \lambda\|\mathbf{W}\boldsymbol{\eta}\|_1. \quad (7.15)$$

Since the covariance of $\boldsymbol{\beta}$ is proportional to $\mathbf{C}^{-1}$ with respect to the posterior distribution (7.13) or (7.14), $\ell_B(\boldsymbol{\eta}, \boldsymbol{\beta})$ is essentially the squared Mahalanobis distance between $\boldsymbol{\eta}$ and $\boldsymbol{\beta}$, plus a weighted ℓ_1 norm of $\boldsymbol{\eta}$ to encourage sparsity. Denote by $\tilde{\boldsymbol{\beta}}$ the optimal decision that minimizes the loss ℓ_B for a given $\boldsymbol{\beta}$, that is, $\tilde{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\eta}} \ell_B(\boldsymbol{\eta}, \boldsymbol{\beta})$. Let $\tilde{\mathbf{S}}$ be the subgradient of $\|\boldsymbol{\eta}\|_1$ at $\tilde{\boldsymbol{\beta}}$. The KKT condition for $\tilde{\boldsymbol{\beta}}$ is

$$\mathbf{C}\tilde{\boldsymbol{\beta}} + \lambda\mathbf{W}\tilde{\mathbf{S}} = \mathbf{C}\boldsymbol{\beta}. \quad (7.16)$$

Since $\boldsymbol{\beta}$ is a random vector in Bayesian inference, the distribution of $\boldsymbol{\beta}$ determines the joint distribution of $\tilde{\boldsymbol{\beta}}$ and $\tilde{\mathbf{S}}$ via the above KKT condition. Represent $(\tilde{\boldsymbol{\beta}}, \tilde{\mathbf{S}})$ by its equivalent form $(\tilde{\boldsymbol{\beta}}_{\tilde{A}}, \tilde{\mathbf{S}}_{\tilde{I}}, \tilde{A})$ in the same way as for $(\hat{\boldsymbol{\beta}}, \mathbf{S})$ in Section 2.

The conditional posterior distribution (7.13) implies that $\mathbf{C}\boldsymbol{\beta} \mid \sigma^2, \mathbf{y} \sim \mathcal{N}_p(\mathbf{C}\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}, n^{-1}\sigma^2\mathbf{C})$. Thus, conditional on $\mathbf{y}$ and σ^2, Equation (7.16) implies that

$$\mathbf{C}\tilde{\boldsymbol{\beta}} + \lambda\mathbf{W}\tilde{\mathbf{S}} - \mathbf{C}\hat{\boldsymbol{\beta}}^{\mathrm{OLS}} \overset{d}{=} \mathbf{U}, \quad (7.17)$$

where $\mathbf{U} \sim \mathcal{N}_p(\mathbf{0}, n^{-1}\sigma^2\mathbf{C})$. One sees that (7.17) is identical to the KKT condition (2.3) with $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ in place of $\boldsymbol{\beta}$. Therefore, the conditional distribution $[\tilde{\boldsymbol{\beta}}_{\tilde{A}}, \tilde{\mathbf{S}}_{\tilde{I}}, \tilde{A} \mid \sigma^2, \mathbf{y}]$, determined by (7.17), is identical to the estimated sampling distribution $\hat{\pi}$ (2.19) under a normal error distribution with $\boldsymbol{\beta}$ estimated by $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$, that is, $\check{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$. Furthermore, $\mathbf{C}\boldsymbol{\beta} \mid \mathbf{y} \sim t_{n-p}(\mathbf{C}\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}, n^{-1}\hat{\sigma}^2\mathbf{C})$ due to (7.14). By a similar reasoning, the conditional distribution $[\tilde{\boldsymbol{\beta}}_{\tilde{A}}, \tilde{\mathbf{S}}_{\tilde{I}}, \tilde{A} \mid \mathbf{y}]$ is the same as $\hat{\pi}$ if $\check{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ and if $f_{\mathbf{U}}$ is estimated by the density of $t_{n-p}(\mathbf{0}, n^{-1}\hat{\sigma}^2\mathbf{C})$. This motivates our proposal to use $t_{n-p}(\mathbf{0}, n^{-1}\sigma^2\mathbf{C})$ as a parametric model for $\mathbf{U}$ and estimate σ^2 from data to construct $\hat{f}_{\mathbf{U}}$. The above discussion also provides a Bayesian justification for sampling from $\hat{\pi}$.

Under this framework, we may define a point estimator $\hat{\boldsymbol{\beta}}^{\mathrm{P}} = (\hat{\beta}_j^{\mathrm{P}})_{1:p}$ by the decision that minimizes the posterior expectation

of the loss $\ell_B(\boldsymbol{\eta}, \boldsymbol{\beta})$,

$$\hat{\boldsymbol{\beta}}^{\mathrm{P}} \triangleq \arg\min_{\boldsymbol{\eta}} \int \ell_B(\boldsymbol{\eta}, \boldsymbol{\beta}) p(\boldsymbol{\beta} \mid \mathbf{y}) d\boldsymbol{\beta}, \tag{7.18}$$

provided that the expectation exists. Although $\hat{\boldsymbol{\beta}}^{\mathrm{P}}$ minimizes the posterior expected loss, its Bayes risk is not well-defined due to our use of an improper prior. To avoid any potential confusion, we call $\hat{\boldsymbol{\beta}}^{\mathrm{P}}$ a posterior point estimator instead of a Bayes estimator. Taking subderivative of $\ell_B(\boldsymbol{\eta}, \boldsymbol{\beta})$ with respect to $\boldsymbol{\eta}$ leads to the following equation to solve for the minimizer $\hat{\boldsymbol{\beta}}^{\mathrm{P}}$:

$$\mathbf{C}\hat{\boldsymbol{\beta}}^{\mathrm{P}} + \lambda \mathbf{W}\mathbf{S}^{\mathrm{P}} = \int \mathbf{C}\boldsymbol{\beta} \cdot p(\boldsymbol{\beta} \mid \mathbf{y}) d\boldsymbol{\beta} = \mathbb{E}(\mathbf{C}\boldsymbol{\beta} \mid \mathbf{y}), \tag{7.19}$$

where $\mathbf{S}^{\mathrm{P}}$ is the subgradient of $\|\boldsymbol{\eta}\|_1$ at $\hat{\boldsymbol{\beta}}^{\mathrm{P}}$. Under the noninformative prior, the posterior mean $\mathbb{E}(\boldsymbol{\beta} \mid \mathbf{y}) = \hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$. In this case, $\mathbb{E}(\mathbf{C}\boldsymbol{\beta} \mid \mathbf{y}) = n^{-1}\mathbf{X}^{\mathrm{T}}\mathbf{y}$ and Equation (7.19) is identical to the KKT condition (2.1) for the Lasso-type estimator $\hat{\boldsymbol{\beta}}$. Therefore, $\hat{\boldsymbol{\beta}}$ can be interpreted as the estimator (7.18) that minimizes the posterior expected loss.

Remark 13. These results provide a Bayesian interpretation of the Lasso-type estimator $\hat{\boldsymbol{\beta}}$ and its sampling distribution. Assume a normal error distribution with a given σ^2 and the noninformative prior. The posterior distribution of the optimal decision, $[\tilde{\boldsymbol{\beta}} \mid \mathbf{y}]$, is identical to the sampling distribution of $\hat{\boldsymbol{\beta}}$ assuming $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ is the true coefficient vector. Therefore, a posterior probability interval for $\tilde{\boldsymbol{\beta}}$, the optimal decision, constructed according to $[\tilde{\boldsymbol{\beta}} \mid \mathbf{y}]$ is the same as the confidence interval constructed according to $\hat{\pi}$ with $\check{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$. Point estimation about $\boldsymbol{\beta}$ also coincides between the Bayesian and the penalized least-squares methods ($\hat{\boldsymbol{\beta}}^{\mathrm{P}} = \hat{\boldsymbol{\beta}}$). Finally, if we set $\lambda = 0$ in the loss (7.15), then the optimal decision $\tilde{\boldsymbol{\beta}}$ is simply $\boldsymbol{\beta}$. In this special case, the aforementioned coincidences become the familiar correspondence between the posterior distribution (7.13) and the sampling distribution of $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$ and that between the posterior mean and $\hat{\boldsymbol{\beta}}^{\mathrm{OLS}}$.

It is worth mentioning that, in a loose sense, this Bayesian interpretation also applies when $p > n$. In this case, the posterior distribution (7.13) does not exist, but $[\mathbf{C}\boldsymbol{\beta} \mid \sigma^2, \mathbf{y}]$ is a well-defined normal distribution in row($\mathbf{X}$). From KKT conditions (7.16) and (7.19), we see that the posterior point estimator $\hat{\boldsymbol{\beta}}^{\mathrm{P}}$ and the posterior distribution $[\tilde{\boldsymbol{\beta}} \mid \mathbf{y}]$ only depend on $\mathbf{C}\boldsymbol{\beta}$. Therefore, they are well-defined and have the same coincidence with the Lasso-type estimator and its sampling distribution.

7.4 Bootstrap Versus Monte Carlo

We have demonstrated that Monte Carlo sampling via estimator augmentation has substantial advantages in approximating tail probabilities and conditional distributions, say $[\hat{\boldsymbol{\beta}}_A \mid \mathcal{A} = A]$, over direct sampling (or bootstrap). The MH Lasso sampler also showed some improvement in efficiency when compared against direct sampling in the low-dimensional setting. Now we discuss some limitations of estimator augmentation relative to bootstrap.

The joint density of the augmented estimator is derived for a given λ, and thus does not take into account the randomness in λ when it is chosen via a data-dependent way, say via cross-validation. Denote by $\hat{\boldsymbol{\beta}}(\mathbf{y}, \hat{\lambda}(\mathbf{y}))$ the Lasso-type estimator when $\lambda = \hat{\lambda}(\mathbf{y})$, where $\hat{\lambda}(\mathbf{y})$ is estimated from the data $\mathbf{y}$. We stress that the density in Theorem 1 or Theorem 2 does not apply to the sampling distribution of $\hat{\boldsymbol{\beta}}(\mathbf{y}, \hat{\lambda}(\mathbf{y}))$ and it is only valid for $\hat{\boldsymbol{\beta}}(\mathbf{y}, \lambda)$ with λ being fixed during the repeated sampling of $\mathbf{y}$. However, the direct sampler (or bootstrap in a similar way) can handle data-dependent λ by adding one additional step to determine $\hat{\lambda}(\mathbf{y}^{(t)})$ after each draw of $\mathbf{y}^{(t)}$ in Routine 1.

Bootstrap and the direct sampler can be parallelized. The IS algorithm (Routine 5) can easily be parallelized as well, since it uses the direct sampler to generate proposals and calculates importance weights independently for each sample. An MCMC algorithm needs a certain number of burn-in iterations before the Markov chain reaches its stationary distribution. It seems that naively running multiple short chains in parallel may impair the overall efficiency due to the computational waist of multiple burn-in iterations. Initialized with one draw from the direct sampler, a Markov chain simulated by Routine 4, however, reaches its equilibrium at the first iteration and thus is suitable for parallel computing. Its efficiency relative to direct sampling when both are parallelized can be calculated as follows.

Suppose our goal is to estimate $\mathbb{E}_\pi[g(\hat{\boldsymbol{\beta}})]$ and assume that $\mathrm{var}_\pi[g(\hat{\boldsymbol{\beta}})] = 1$ without loss of generality. Assume that the time to run one iteration of the direct sampler allows for running m iterations of an MCMC algorithm. Suppose that we have access to K computing nodes and the available computing time from each node allows for the simulation of $(1 + N_1)$ samples from the direct sampler, where N_1 may be small. Thus, on a single node we can run $N_2 = mN_1$ MCMC iterations plus an initial draw from the direct sampler in the same amount of time. In other words, we can run Routine 4 for $1 + N_2$ iterations to draw $\boldsymbol{\beta}^{(t)}$ for $t = 1, \ldots, 1 + N_2$. Note that this Markov chain reaches equilibrium from $t = 1$. Let $\rho_t = \mathrm{cor}(g(\boldsymbol{\beta}^{(1)}), g(\boldsymbol{\beta}^{(t+1)}))$ and

$$\psi(N) = 1 + 2\sum_{t=1}^{N-1} \left(1 - \frac{t}{N}\right) \rho_t$$

for an integer $N \geq 1$. Then, we have

$$\mathrm{var}\left[\frac{1}{N_2+1}\sum_{t=1}^{N_2+1} g(\boldsymbol{\beta}^{(t)})\right] = \frac{1}{N_2+1}\psi(N_2+1) \triangleq V_2(N_2+1).$$

Denote by $V_1(N) = 1/N$ the variance in estimating g by the mean of an iid sample of size N, and let

$$\gamma = \lim_{N\to\infty} \frac{V_1(N)}{V_2(mN)} = \frac{m}{\psi(\infty)}.$$

The efficiency of Routine 4 relative to direct sampling is

$$\begin{aligned}\frac{V_1(N_1+1)}{V_2(N_2+1)} &= \frac{mN_1+1}{N_1+1}\frac{1}{\psi(N_2+1)} \\ &> \frac{N_1}{N_1+1}\frac{m}{\psi(N_2+1)} \geq \frac{N_1}{N_1+1}\gamma,\end{aligned}$$

where we have assumed that $\psi(N_2 + 1) \leq \psi(\infty)$ for the last inequality. This assumption holds if $\psi(N)$ is nondecreasing in N. This derivation shows that Routine 4 will be more efficient than the direct sampler on each computing node if $N_1 \geq 1/(\gamma - 1)$, which can be as small as 1 when $\gamma > 2$. We have observed two decay patterns of the autocorrelation ρ_t of the MLS in

the simulation study in Section 3.6. For some components of $\hat{\boldsymbol{\beta}}$, ρ_t is always positive before it decays to zero, in which case $\psi(N)$ is obviously nondecreasing. For other components, ρ_t first decreases monotonely to zero and then shows small fluctuations around zero. In the second case, we empirically observed that $\psi(N)$ is nondecreasing as well. The efficiency comparison in Table 2, with N_1 and $N_2 = mN_1$ both large, suggests that for most functions estimated there, $\gamma \in (2, 3)$ for datasets A and B and $\gamma \in (1.2, 1.6)$ for the other two datasets. Therefore, as long as we need to run a few iterations of the direct sampler on each node, parallelizing Routine 4 can bring computational gain. Of course, if the number of computing nodes K is so large that only one draw is needed from each node, direct sampling or bootstrap will be a better choice.

7.5 Concluding Remarks

Using the density of an augmented estimator, this article develops MCMC and IS methods to approximate sampling distributions in ℓ_1-penalized linear regression. This approach is clearly different from existing methods based on resampling or asymptotic approximation. The numerical results have already demonstrated the substantial gain in efficiency and the great flexibility offered by this approach. These results are mostly for a proof of principle, and there is room for further development of more efficient Monte Carlo algorithms based on the densities derived in this article.

In principle, the idea of estimator augmentation can be applied to the use of concave penalties in linear regression (Frank and Friedman 1993; Fan and Li 2001; Friedman, Hastie, and Tibshirani 2008; Zhang 2010) for studying the sampling distribution. However, there are at least two additional technical difficulties for the high-dimensional setting. First, we need to find conditions for the uniqueness of a concave-penalized estimator to construct a bijection between $\mathbf{U}$ and the augmented estimator. Second, the constraint in (4.4) will become nonlinear in general, even for a fixed $\mathbf{s}_A$, when a concave penalty is used, which means that the sample space is composed of a finite number of manifolds. Another future direction is to investigate theoretically and empirically the finite-sample performance in variable selection by the Lasso sampler which may take into account the uncertainty in parameter estimation in a coherent way.

8. PROOFS

Let $n \geq 1$ and $p \geq 2$ throughout this section.

8.1 Proof of Theorem 3

Lemma 11. Let $\mathbf{Z} \in \mathbb{R}^p$ be a random vector, $K \in \mathscr{R}^p$, and $\mathbf{Z}_K$ be the truncation of $\mathbf{Z}$ to K such that $P(\mathbf{Z}_K \in B) = P(\mathbf{Z} \in B \mid \mathbf{Z} \in K)$ for $B \in \mathscr{R}^p$. If $P(\mathbf{Z} \in K) \geq 1 - \alpha > 0$, then

$$\sup_{B \in \mathscr{R}^p} |P(\mathbf{Z}_K \in B) - P(\mathbf{Z} \in B)| \leq \alpha.$$

Proof. For $B \in \mathscr{R}^p$, $P(\mathbf{Z} \in B \cap K) = P(\mathbf{Z}_K \in B)P(\mathbf{Z} \in K)$ and thus

$$0 \leq P(\mathbf{Z}_K \in B) - P(\mathbf{Z} \in B \cap K) = P(\mathbf{Z}_K \in B) \cdot P(\mathbf{Z} \in K^c) \leq \alpha.$$

On the other hand,

$$0 \leq P(\mathbf{Z} \in B) - P(\mathbf{Z} \in B \cap K) = P(\mathbf{Z} \in B \cap K^c) \leq \alpha.$$

Therefore, $|P(\mathbf{Z}_K \in B) - P(\mathbf{Z} \in B)| \leq \alpha$ for any B and the conclusion follows. □

Lemma 12. Assume that $\check{\boldsymbol{\beta}}$ satisfies conditions (i) and (ii) in Lemma 7, and let (6.9) be satisfied. Then, $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{u}) = V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{u})$ for any $\mathbf{u} \in \mathbb{R}^p$ if $\|\boldsymbol{\delta}\|_\infty \leq M_1$.

Proof. The assumptions on $\check{\boldsymbol{\beta}}$ imply that $\mathcal{A}(\check{\boldsymbol{\beta}}) = A_0$ and $\operatorname{sgn}(\check{\beta}_j) = \operatorname{sgn}(\beta_{0j})$ for all $j \in A_0$. By the definition of V (6.1) it then suffices to show that

$$\left|\beta_{0j} + r_n^{-1}\delta_j\right| - |\beta_{0j}| = \left|\check{\beta}_j + r_n^{-1}\delta_j\right| - |\check{\beta}_j| \tag{8.1}$$

for all $j \in A_0$. By the definition of η in (6.8), for $j \in A_0$

$$\eta|\beta_{0j}| \geq |\check{\beta}_j - \beta_{0j}| \geq |\beta_{0j}| - |\check{\beta}_j|$$

and therefore

$$|\check{\beta}_j| \geq (1 - \eta)|\beta_{0j}| > M_1/r_n \geq |r_n^{-1}\delta_j|,$$

where we have used (6.9) and that $\|\boldsymbol{\delta}\|_\infty \leq M_1$. Consequently, for $j \in A_0$ we have

$$|\check{\beta}_j + r_n^{-1}\delta_j| - |\check{\beta}_j| = \operatorname{sgn}(\check{\beta}_j)\delta_j/r_n.$$

On the other hand, by (6.9) and $\eta \in [0, 1)$, $|\beta_{0j}| > M_1/r_n \geq |r_n^{-1}\delta_j|$ for $j \in A_0$ and thus

$$|\beta_{0j} + r_n^{-1}\delta_j| - |\beta_{0j}| = \operatorname{sgn}(\beta_{0j})\delta_j/r_n.$$

Now (8.1) follows since $\operatorname{sgn}(\check{\beta}_j) = \operatorname{sgn}(\beta_{0j})$ for all $j \in A_0$. □

Proof of Lemma 7. Define $\tilde{\boldsymbol{\delta}} = \arg\min_{\boldsymbol{\delta}} V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}^*)$, which follows the same distribution as $\hat{\boldsymbol{\delta}}$ (6.2). This is because $\mathbf{U}^* \overset{d}{=} \mathbf{U}$ by fixing $\hat{\sigma} = \sigma$ and $V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{u})$ has a unique minimizer for any $\mathbf{u}$ if the columns of $\mathbf{X}$ are in general position (Lemma 1). Consequently,

$$\nu[\hat{\boldsymbol{\delta}} \mid \|\hat{\boldsymbol{\delta}}\|_\infty < M_1] = \nu[\tilde{\boldsymbol{\delta}} \mid \|\tilde{\boldsymbol{\delta}}\|_\infty < M_1].$$

Let $\mathcal{K} = \{\boldsymbol{\delta} \in \mathbb{R}^p : \|\boldsymbol{\delta}\|_\infty < M_1\}$. According to Lemma 12, $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{U}^*) = V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}^*)$ for all $\boldsymbol{\delta} \in \mathcal{K}$. As the unique minimizer of $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{U}^*)$ (6.3), $\|\boldsymbol{\delta}^*\|_\infty < M_1$ implies that $\boldsymbol{\delta}^*$ is also a local minimizer of $V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}^*)$. Since $V(\boldsymbol{\delta}; \boldsymbol{\beta}_0, \mathbf{U}^*)$ is convex in $\boldsymbol{\delta}$ and has only a unique minimizer $\tilde{\boldsymbol{\delta}}$, we must have $\boldsymbol{\delta}^* = \tilde{\boldsymbol{\delta}}$ and $\|\tilde{\boldsymbol{\delta}}\|_\infty < M_1$. Furthermore, using the same argument in the other direction, one can show that $\|\tilde{\boldsymbol{\delta}}\|_\infty < M_1$ implies $\|\boldsymbol{\delta}^*\|_\infty < M_1$, and thus $\{\|\tilde{\boldsymbol{\delta}}\|_\infty < M_1\}$ is equivalent to $\{\|\boldsymbol{\delta}^*\|_\infty < M_1\}$. This completes the proof.

Proof of Theorem 3. Let E_1 be the event that $\check{\boldsymbol{\beta}}$ satisfies conditions (i) and (ii) in Lemma 7 and $E_2 = \{\|\hat{\boldsymbol{\delta}}\|_\infty < M_1\}$. We first show that (6.12) holds on E_1. Obviously, (6.9) holds because of (6.11). The argument in the proof of Lemma 7 implies that, on event E_1,

$$\begin{aligned} P(\|\boldsymbol{\delta}^*\|_\infty < M_1 \mid \check{\boldsymbol{\beta}}) &= P(\|\tilde{\boldsymbol{\delta}}\|_\infty < M_1) \\ &= P(E_2) \geq 1 - \alpha_1, \end{aligned} \tag{8.2}$$

where the second equality is due to $\tilde{\boldsymbol{\delta}} \overset{d}{=} \hat{\boldsymbol{\delta}}$. Let $\boldsymbol{\delta}^*_{\mathcal{K}}$ and $\hat{\boldsymbol{\delta}}_{\mathcal{K}}$ be the respective truncations of $\boldsymbol{\delta}^*$ and $\hat{\boldsymbol{\delta}}$ to $\mathcal{K}$. Lemma 7 implies $\nu[\boldsymbol{\delta}^*_{\mathcal{K}} \mid$

$\check{\boldsymbol{\beta}}] = \nu[\hat{\boldsymbol{\delta}}_{\mathcal{K}}]$ on event E_1. A direct consequence is that on E_1, $P(\boldsymbol{\delta}^*_{\mathcal{K}} \in B \mid \check{\boldsymbol{\beta}}) = P(\hat{\boldsymbol{\delta}}_{\mathcal{K}} \in B)$ for any $B \in \mathscr{R}^p$ and therefore,

$$\sup_{B \in \mathscr{R}^p} |P(\boldsymbol{\delta}^* \in B \mid \check{\boldsymbol{\beta}}) - P(\hat{\boldsymbol{\delta}} \in B)| \le 2\alpha_1$$

by Lemma 11 and (8.2).

Next we find a lower bound for $P(E_1)$. Since $\inf_{A_0} |\beta_{0j}| > 2\tau$ (6.11), on event $\mathcal{E}$, we have $\mathcal{A}(\check{\boldsymbol{\beta}}) = A_0$, according to Lemma 6. By construction $\check{\boldsymbol{\beta}}_{A_0} = \hat{\boldsymbol{\beta}}_{A_0}$ (6.10) and consequently

$$\eta = \sup_{j \in A_0} \frac{|\hat{\beta}_j - \beta_{0j}|}{|\beta_{0j}|} \le \frac{\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_2}{\inf_{A_0} |\beta_{0j}|} < \frac{1}{2} \tag{8.3}$$

again on $\mathcal{E}$. Thus, $P(E_1) \ge P(\mathcal{E}) \ge 1 - 2p^{1-u}$.

8.2 Proof of Theorem 4

Lemma 13. Let $\boldsymbol{\gamma}$ be any minimizer of $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{u})$ for $\check{\boldsymbol{\beta}} \in \mathbb{R}^p$ and $\mathbf{u} \in \mathbb{R}^p$. For any $\boldsymbol{\Delta} \in \mathbb{R}^p$, we have

$$V(\boldsymbol{\gamma} + \boldsymbol{\Delta}; \check{\boldsymbol{\beta}}, \mathbf{u}) \ge V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, \mathbf{u}) + \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta}. \tag{8.4}$$

Proof. Let $\mathbf{b} = r_n^{-1} \boldsymbol{\gamma} + \check{\boldsymbol{\beta}} = (b_j)_{1:p}$ and $\boldsymbol{\Delta} = (\Delta_j)_{1:p}$. Direct calculations give

$$V(\boldsymbol{\gamma} + \boldsymbol{\Delta}; \check{\boldsymbol{\beta}}, \mathbf{u}) - V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, \mathbf{u}) = \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} (\boldsymbol{\Delta} + 2\boldsymbol{\gamma}) - \frac{n}{r_n} \boldsymbol{\Delta}^\mathsf{T} \mathbf{u} + n\lambda \sum_{j=1}^p w_j (|b_j + r_n^{-1} \Delta_j| - |b_j|).$$

The KKT condition for $\boldsymbol{\gamma}$ to minimize $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{u})$ is

$$\mathbf{C}(\mathbf{b} - \check{\boldsymbol{\beta}}) + \lambda \mathbf{W} \mathbf{s} - \mathbf{u} = \mathbf{0}, \tag{8.5}$$

where $\mathbf{s} = (s_j)_{1:p}$ is the subgradient of $\|\boldsymbol{\beta}\|_1$ at $\mathbf{b}$. By the definition of a subgradient,

$$|b_j + r_n^{-1} \Delta_j| - |b_j| \ge s_j r_n^{-1} \Delta_j \tag{8.6}$$

for all $j = 1, \ldots, p$. Now we have

$$\begin{aligned} &V(\boldsymbol{\gamma} + \boldsymbol{\Delta}; \check{\boldsymbol{\beta}}, \mathbf{u}) - V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, \mathbf{u}) \\ &\ge \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta} + \frac{n}{r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\gamma} - \frac{n}{r_n} \boldsymbol{\Delta}^\mathsf{T} \mathbf{u} + \frac{n\lambda}{r_n} \sum_{j=1}^p w_j s_j \Delta_j \\ &= \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta} + \frac{n}{r_n} \boldsymbol{\Delta}^\mathsf{T} [\mathbf{C}(\mathbf{b} - \check{\boldsymbol{\beta}}) - \mathbf{u} + \lambda \mathbf{W} \mathbf{s}] \\ &= \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta}, \end{aligned}$$

where we have used (8.6) and (8.5). □

Lemma 14. Assume that $|\mathcal{A}(\check{\boldsymbol{\beta}})| \le q$. Let $c > 0$, $\mathbf{u} = (u_j)_{1:p} \in \mathbb{R}^p$, and $\boldsymbol{\gamma}$ be any minimizer of $V(\boldsymbol{\delta}; \check{\boldsymbol{\beta}}, \mathbf{u})$. If Assumption RE$(q, 3)$ is satisfied and $|u_j| \le w_j \lambda / 2$ for all $j = 1, \ldots, p$, then

$$|V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, c\mathbf{u}) - V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, \mathbf{u})| \le \frac{8 w_{\max} |1 - c|}{w_{\min} \kappa^2(q, 3)} n\lambda^2 \sum_{A_0} w_j^2. \tag{8.7}$$

Proof. Let $\mathbf{b} = r_n^{-1} \boldsymbol{\gamma} + \check{\boldsymbol{\beta}}$. Direct calculations give

$$\begin{aligned} |V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, c\mathbf{u}) - V(\boldsymbol{\gamma}; \check{\boldsymbol{\beta}}, \mathbf{u})| &= n \left| (1 - c) \mathbf{u}^\mathsf{T} (\mathbf{b} - \check{\boldsymbol{\beta}}) \right| \\ &\le n |1 - c| \cdot \|\mathbf{u}\|_\infty \|\mathbf{b} - \check{\boldsymbol{\beta}}\|_1. \end{aligned}$$

It is seen from (8.5) that $\mathbf{b}$ is a minimizer of the loss (1.2) if $\check{\boldsymbol{\beta}}$ is the true coefficient vector and if $\mathbf{X}^\mathsf{T} \varepsilon / n = \mathbf{u}$. Inequality (6.6) in Lemma 5 applied under these assumptions leads to

$$\|\mathbf{b} - \check{\boldsymbol{\beta}}\|_1 \le \frac{16\lambda}{\kappa^2(q, 3)} \sum_{A_0} \frac{w_j^2}{w_{\min}}$$

if $|u_j| \le w_j \lambda / 2$ for all j. Moreover, $\|\mathbf{u}\|_\infty \le w_{\max} \lambda / 2$ and hence (8.7) follows. □

Proof of Lemma 8. To simplify notation, let $\hat{c} = \sigma / \hat{\sigma}$ and

$$h = \frac{8 w_{\max} \zeta}{w_{\min} \kappa^2(q, 3)} n\lambda^2 \sum_{A_0} w_j^2.$$

If $|U_j^*| \le (1 - \zeta) w_j \lambda / 2$, then $|U_j^*| \le w_j \lambda / 2$ and by (6.17) $|\hat{c} U_j^*| \le w_j \lambda / 2$. Lemma 14 with (6.17) implies

$$\begin{aligned} |V(\boldsymbol{\delta}^*; \check{\boldsymbol{\beta}}, \hat{c}\mathbf{U}^*) - V(\boldsymbol{\delta}^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*)| &\le h, \\ |V(\boldsymbol{\delta}_0^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*) - V(\boldsymbol{\delta}_0^*; \check{\boldsymbol{\beta}}, \hat{c}\mathbf{U}^*)| &\le h. \end{aligned}$$

Let $\boldsymbol{\Delta} = \boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*$. Now we have

$$\begin{aligned} V(\boldsymbol{\delta}^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*) &\ge V(\boldsymbol{\delta}^*; \check{\boldsymbol{\beta}}, \hat{c}\mathbf{U}^*) - h \\ &\ge V(\boldsymbol{\delta}_0^*; \check{\boldsymbol{\beta}}, \hat{c}\mathbf{U}^*) + \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta} - h \\ &\ge V(\boldsymbol{\delta}_0^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*) + \frac{n}{2r_n^2} \boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta} - 2h, \end{aligned}$$

where the second inequality is due to Lemma 13. Finally, since $V(\boldsymbol{\delta}_0^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*) \ge V(\boldsymbol{\delta}^*; \check{\boldsymbol{\beta}}, \mathbf{U}^*)$ by definition (6.3), $\boldsymbol{\Delta}^\mathsf{T} \mathbf{C} \boldsymbol{\Delta} \le 4 r_n^2 h / n$ which coincides with (6.18).

Proof of Theorem 4. Recall that $\mathcal{E}$ is the event $\cap_{j=1}^p \{|U_j| \le w_j \lambda / 2\}$. Since the distribution of $\boldsymbol{\delta}_0^*$ does not depend on $\hat{\sigma}$ and is identical to the distribution of $\boldsymbol{\delta}^*$ when $\hat{\sigma}$ is fixed to the true noise level σ, (6.21) follows immediately from (6.12) which holds on $\mathcal{E}$.

Let E_3 be the event in (6.17) and $\mathcal{E}^*$ be the event that $\cap_{j=1}^p \{|U_j^*| \le (1 - \zeta) w_j \lambda / 2\}$. By Lemma 6, on $\mathcal{E}$ we have $\mathcal{A}(\check{\boldsymbol{\beta}}) = A_0$, $|\mathcal{A}(\check{\boldsymbol{\beta}})| = q_0 \le q$ and by (8.3)

$$\inf_{A_0} |\check{\beta}_j| \ge \frac{1}{2} \inf_{A_0} |\beta_{0j}| > \tau. \tag{8.8}$$

Therefore, all the assumptions on $\check{\boldsymbol{\beta}}$ and $\hat{\sigma}$ in Lemma 8 are satisfied on $\mathcal{E} \cap E_3$, which happens with probability at least $1 - (\alpha_2 + 2p^{1-u})$. Moreover, the conditional probability of (6.18) given $(\check{\boldsymbol{\beta}}, \hat{\sigma})$ is at least

$$P(\mathcal{E}^* \mid \hat{\sigma}) \ge 1 - 2p^{1-u},$$

by choosing $\lambda \ge (\hat{\sigma} / \sigma) \lambda_0 / (1 - \zeta)$. For the lower bound of the above probability, see (6.4) in Lemma 5 with $(1 - \zeta) w_j$ in place of w_j and $\hat{\sigma}$ in place of σ. As $\hat{\sigma} / \sigma \le 1 + \zeta$ on E_3, it suffices to choose λ as in (6.20). What remains is to show that $\boldsymbol{\Delta} = \boldsymbol{\delta}^* - \boldsymbol{\delta}_0^*$ is M_2-sparse on the event $\mathcal{E}^*$. Then, (6.22) follows from (6.18) and the definition of $\phi_{\min}(M_2)$ (6.19). Regarding $\check{\boldsymbol{\beta}}$ and $\hat{\sigma}$ as the true parameters, Lemma 6 with (8.8) implies that $\mathcal{A}(\check{\boldsymbol{\beta}}) \subseteq \mathcal{A}(\boldsymbol{\beta}^*)$

on $\mathcal{E}^*$ and therefore, $|\mathcal{A}(\boldsymbol{\delta}^*)| \le |\mathcal{A}(\boldsymbol{\beta}^*)| \le M_2/2$ by (6.5). Since $\mathcal{E}^*$ with (6.17) implies $|(\sigma/\hat{\sigma})U_j^*| \le w_j\lambda/2$ for all j, by a similar reasoning we also have $|\mathcal{A}(\boldsymbol{\delta}_0^*)| \le M_2/2$ and thus $|\mathcal{A}(\boldsymbol{\Delta})| \le M_2$ on $\mathcal{E}^*$.

APPENDIX

Recall that $A^\dagger = A \setminus \{j\}$ in proposal (P3) and $A^\dagger = A \cup \{j\}$ in (P4). Let $B = A \cap A^\dagger$. For both proposals,

$$\frac{\det \mathbf{C}_{A^\dagger A^\dagger}}{\det \mathbf{C}_{AA}} = \left(\mathbf{C}_{jj} - \mathbf{C}_{jB}\mathbf{C}_{BB}^{-1}\mathbf{C}_{Bj}\right)^{|A^\dagger|-|A|} \triangleq (r_{\text{det}})^{|A^\dagger|-|A|}. \tag{A.1}$$

Suppose that the matrix $\mathbf{C}_{AA}^{-1}$ is given.

When (P3) is proposed, let $k(j) \in \{1, \ldots, |A|\}$ index the position of j in the set A and d_k be the kth diagonal element of $\mathbf{C}_{AA}^{-1}$. Then, $d_{k(j)} = 1/r_{\text{det}}$ and thus the ratio (A.1) is immediately obtained. If this proposal is rejected, no further computation is necessary. If it is accepted, $\mathbf{C}_{A^\dagger A^\dagger}^{-1}$ can be obtained after a reverse sweeping of $(-\mathbf{C}_{AA}^{-1})$ on position $k(j)$. When (P4) is proposed, $r_{\text{det}} = \mathbf{C}_{jj} - \mathbf{C}_{jA}\mathbf{C}_{AA}^{-1}\mathbf{C}_{Aj}$ and thus the ratio (A.1) can be readily calculated. Again, if the proposal is rejected, no further computation is needed. If it is accepted, add j to the last position in the set $A^\dagger$ and then sweep the matrix

$$\begin{pmatrix} -\mathbf{C}_{AA}^{-1} & \mathbf{C}_{AA}^{-1}\mathbf{C}_{Aj} \\ \mathbf{C}_{jA}\mathbf{C}_{AA}^{-1} & r_{\text{det}} \end{pmatrix}$$

on the last position to obtain $-\mathbf{C}_{A^\dagger A^\dagger}^{-1}$. It is seen that for both proposals, the ratio (A.1) can be calculated easily and sweeping on a single position is all we need to update $\mathbf{C}_{AA}^{-1}$.

[Received October 2012. Revised June 2014.]

REFERENCES

Bach, F. (2008), "Bolasso: Model Consistent Lasso Estimation Through the Bootstrap," in *Proceedings of the 25th International Conference on Machine Learning*, New York: Association for Computing Machinery, pp. 33–40. [1495]

Bickel, P., Ritov, Y., and Tsybakov, A. (2009), "Simultaneous Analysis of Lasso and Dantzig Selector," *The Annals of Statistics*, 37, 1705–1732. [1504,1508]

Chatterjee, A., and Lahiri, S. (2010), "Asymptotic Properties of the Residual Bootstrap for Lasso Estimators," *Proceedings of the American Mathematical Society*, 138, 4497–4509. [1495]

——— (2011), "Boostrapping Lasso Estimators," *Journal of the American Statistical Association*, 106, 608–625. [1495,1510]

——— (2013), "Rates of Convergence of the Adaptive Lasso Estimators to the Oracle Distribution and Higher Order Refinements by the Bootstrap," *The Annals of Statistics*, 41, 1232–1259. [1510]

Chen, S., Donoho, D. L., and Saunders, M. (1999), "Atomic Decomposition by Basis Pursuit," *SIAM Journal of Scientific Computing*, 20, 33–61. [1495]

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), "Least Angle Regression," *The Annals of Statistics*, 32, 407–499. [1495]

Fan, J., and Li, R. (2001), "Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties," *Journal of the American Statistical Association*, 96, 1348–1360. [1495,1502,1514]

Frank, I., and Friedman, J. (1993), "A Statistical View of Some Chemometrics Regression Tools" (with discussion), *Technometrics*, 35, 109–148. [1514]

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), "Pathwise Coordinate Optimization," *The Annals of Applied Statistics*, 1, 302–332. [1495,1507]

Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation with the Graphical Lasso," *Biostatistics*, 9, 432–441. [1514]

Green, P. (1995), "Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination," *Biometrika*, 82, 711–732. [1501]

Javanmard, A., and Montanari, A. (2013a), "Confidence Intervals and Hypothesis Testing for High-Dimensional Regression," *arXiv*, 1306.3171. [1495]

——— (2013b), "Hypothesis Testing in High-Dimensional Regression Under the Gaussian Random Design Model: Asymptotic Theory," *arXiv*, 1301.4240. [1495]

Knight, K., and Fu, W. (2000), "Asymptotics for Lasso-Type Estimators," *The Annals of Statistics*, 28, 1356–1378. [1495,1501,1510]

Lockhart, R., Taylor, J., Tibshirani, R., and Tibshirani, R. (2014), "A Significance Test for the Lasso," *The Annals of Statistics*, 42, 413–468. [1495]

Lounici, K., Pontil, M., van de Geer, S., and Tsybakov, A. (2011), "Oracle Inequalities and Optimal Inference Under Group Sparsity," *The Annals of Statistics*, 39, 2164–2204. [1508]

Meinshausen, N., and Bühlmann, P. (2006), "High-Dimensional Graphs and Variable Selection with the Lasso," *The Annals of Statistics*, 34, 1436–1462. [1504,1510,1512]

Meinshausen, N., and Bühlmann, P. (2010), "Stability Selection" (with discussion), *Journal of the Royal Statistical Society,* Series B, 72, 417–473. [1495]

Meinshausen, N., and Yu, B. (2009), "Lasso-Type Recovery of Sparse Representations for High-Dimensional Data," *The Annals of Statistics*, 37, 246–270. [1511]

Minnier, J., Tian, L., and Cai, T. (2011), "A Perturbation Method for Inference on Regularized Regression Estimates," *Journal of the American Statistical Association*, 106, 1371–1382. [1495]

Osborne, M., Presnell, B., and Turlach, B. (2000), "A New Approach to Variable Selection in Least Squares Problems," *IMA Journal of Numerical Analysis*, 20, 389–404. [1495]

Pötscher, B. M., and Schneider, U. (2009), "On the Distribution of the Adaptive LASSO Estimator," *Journal of Statistical Planning and Inference*, 139, 2775–2790. [1495]

——— (2010), "Confidence Sets Based on Penalized Maximum Likelihood Estimators in Gaussian Regression," *Electronic Journal of Statistics*, 4, 334–360. [1495]

Sun, T., and Zhang, C. (2012), "Scaled Sparse Linear Regression," *Biometrika*, 99, 879–898. [1510]

Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society,* Series B, 58, 267–288. [1495]

Tibshirani, R. (2013), "The Lasso Problem and Uniqueness," *Electronic Journal of Statistics*, 7, 1456–1490. [1496,1505]

van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2013), "Confidence Regions and Tests for High-Dimensional Models," *arXiv*, 1303.0518. [1495]

Wainwright, M. (2009), "Sharp Thresholds for High-Dimensional and Noisy Sparsity Recovery Using ℓ_1-Constrained Quadratic Programming (Lasso)," *IEEE Transactions on Information Theory*, 55, 2183–2202. [1512]

Wu, T., and Lange, K. (2008), "Coordinate Descent Procedures for Lasso Penalized Regression," *The Annals of Applied Statistics*, 2, 224–244. [1495,1507]

Yuan, M., and Lin, Y. (2006), "Model Selection and Estimation in Regression With Grouped Variables," *Journal of the Royal Statistical Society,* Series B, 68, 49–67. [1508]

Zhang, C. (2010), "Nearly Unbiased Variable Selection Under Minimax Concave Penalty," *The Annals of Statistics*, 38, 894–942. [1514]

Zhang, C., and Huang, J. (2008), "The Sparsity and Bias of the LASSO Selection in High-Dimensional Linear Regression," *The Annals of Statistics*, 36, 1567–1594. [1504]

Zhang, C., and Zhang, S. (2014), "Confidence Intervals for Low-Dimensional Parameters in High-Dimensional Linear Models," *Journal of the Royal Statistical Society,* Series B, 76, 217–242. [1495]

Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," *Journal of Machine Learning Research*, 7, 2541–2563. [1504,1510,1512]

Zou, H. (2006), "The Adaptive Lasso and Its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429. [1495,1502,1510,1512]

Zou, H., and Li, R. (2008), "One-Step Sparse Estimates in Nonconcave Penalized Likelihood Models," *The Annals of Statistics*, 36, 1509–1533. [1495]