

NONPENALIZED VARIABLE SELECTION IN HIGH-DIMENSIONAL LINEAR MODEL SETTINGS VIA GENERALIZED FIDUCIAL INFERENCE

BY JONATHAN P. WILLIAMS AND JAN HANNIG

University of North Carolina at Chapel Hill

Standard penalized methods of variable selection and parameter estimation rely on the magnitude of coefficient estimates to decide which variables to include in the final model. However, coefficient estimates are unreliable when the design matrix is collinear. To overcome this challenge, an entirely new perspective on variable selection is presented within a generalized fiducial inference framework. This new procedure is able to effectively account for linear dependencies among subsets of covariates in a high-dimensional setting where p can grow almost exponentially in n , as well as in the classical setting where $p \leq n$. It is shown that the procedure very naturally assigns small probabilities to subsets of covariates which include redundancies by way of explicit L_0 minimization. Furthermore, with a typical sparsity assumption, it is shown that the proposed method is consistent in the sense that the probability of the true sparse subset of covariates converges in probability to 1 as $n \rightarrow \infty$, or as $n \rightarrow \infty$ and $p \rightarrow \infty$. Very reasonable conditions are needed, and little restriction is placed on the class of possible subsets of covariates to achieve this consistency result.

1. Introduction. A strategy for developing variable selection procedures with desirable consistency properties entails exploiting some distinguishing property of the theoretical true data generating model. For example, standard penalized methods of variable selection within a linear model framework such as LASSO of Tibshirani [21], SCAD of Fan and Li [9] and the Dantzig Selector of Candès and Tao [8] rely on the magnitude of the coefficients in the true data generating model being relatively larger than those of the other coefficients. Johnson and Rossell [13] use this property to construct nonlocal prior densities over all subsets of covariates (see also Rossell and Telesca [19] and Shin, Bhattacharya and Johnson [20]). The defining property of their nonlocal density is that it takes the value of zero for subsets containing a covariate with a zero-valued coefficient.

We propose a more desirable way for eliminating redundancies from the sample space of candidate subsets which does not explicitly rely on coefficient magnitudes. That is, any candidate true model should be nonredundant in the sense that it

Received February 2018.

MSC2010 subject classifications. 62J05, 62F12, 62A01.

Key words and phrases. Best subset selection, high-dimensional regression, L_0 minimization, feature selection.

contains the minimal amount of information necessary for explaining and/or predicting the observed data. One such criterion to exploit this nonredundancy property is that the only subsets with nonzero posterior probability should be those which cannot be predicted to some chosen precision by a subset of fewer covariates. Such a criterion requires constructing a probability distribution on the space of candidate models, which is consistent with a Bayesian or fiducial variable selection paradigm. The literature on high-dimensional linear models is vast, but we hope to contribute to it by using this setting to build a foundation for a fresh perspective on variable selection.

Recent work in the Bayesian high-dimensional linear model setting includes Ročková and George [18] who develop methods for separable and nonseparable spike-and-slab penalized estimation, the credible set approach of Bondell and Reich [6], and Narisetty and He [16] who propose a method based on shrinking and diffusing coefficient priors in which the variance of the priors are sample size dependent. Lai, Hannig and Lee [14] layout framework for penalized estimation within a GFI approach.

Ghosh and Ghattas [10] provide insights into complications in Bayesian variable selection. Namely, the size of the sample space (2^p) is often too large to compute all model probabilities, and even typically larger than can reasonably be sampled by Markov chain Monte Carlo (MCMC) methods. Thus, the nonlocal prior approach of Johnson and Rossell [13] can achieve asymptotic consistency (where other approaches can only achieve asymptotic pairwise consistency) because it is able to effectively eliminate a large enough portion of the 2^p subsets from the sample space. To illustrate this point, consider the following simple example. Let

$$(1) \quad Y \sim N_n(\beta_1 \cdot x^{(1)} + \dots + \beta_p \cdot x^{(p)}, \sigma^2 I_n),$$

where $\beta_j \in \mathbb{R}$ and $x^{(j)} \in \mathbb{R}^n$ for $j \in \{1, \dots, p\}$, and $\sigma > 0$. Further, suppose that the true but unknown values of $(\beta_1, \beta_2, \beta_3, \dots, \beta_p)'$ are $(b_1, b_2, 0, \dots, 0)'$. Within the nonlocal prior framework, the only subsets with nonnegligible posterior probability are contained in the set $\{\{x^{(1)}\}, \{x^{(2)}\}, \{x^{(1)}, x^{(2)}\}\}$.

When viewed as a prior density on the coefficients, nonlocal priors assign zero prior density to the true parameter value when the true parameter value is zero. From a Bayesian perspective, this is philosophically problematic, but very insightful for consistency of model selection. The insight lends itself to the question: What other properties might the true model have which can be exploited to develop a statistical procedure with the ability to effectively eliminate subsets from the sample space?

In addressing this question, we build our proposed methods from the idea that any candidate true model, as determined by the actual nonzero parameter values, should be nonredundant in the sense that it contains the minimal amount of information necessary for explaining and/or predicting the observed data. We denote such subsets of the parameter space as ε -admissible, and define them precisely

in Definition 2.1. Then, using the above nonlocal prior example, the entire model space of $\{x^{(1)}, x^{(2)}, x^{(3)}\}$ for instance, is not ε -admissible because it can be perfectly predicted by the smaller subset $\{x^{(1)}, x^{(2)}\}$.

To further illustrate the intuition behind our proposal, consider an example where $x^{(2)}$ is highly collinear with all of $x^{(3)}, \dots, x^{(p)}$ but is *not* correlated with $x^{(1)}$, and where the true values of $(\beta_1, \beta_2, \beta_3, \dots, \beta_p)'$ are $(b_1, b_2, b_3, \dots, b_p)'$ with $b_j \neq 0$ for all $j \in \{1, \dots, p\}$. In this case, assuming strong enough collinearity, $\exists c \in \mathbb{R}$ with $c \cdot x^{(2)} \approx b_2 \cdot x^{(2)} + \dots + b_p \cdot x^{(p)}$, that is,

$$\|(b_1 \cdot x^{(1)} + \dots + b_p \cdot x^{(p)}) - (b_1 \cdot x^{(1)} + c \cdot x^{(2)})\| < \varepsilon,$$

where $\varepsilon > 0$ is some desired precision. Thus, for much of the parameter space the subset $\{x^{(1)}, \dots, x^{(p)}\}$ is not ε -admissible, but would be assigned nonzero posterior probability in the nonlocal prior framework.

We construct a posterior-like probability distribution over all subsets, which assigns negligible probability to elements that are not ε -admissible. In constructing the posterior-like probability distribution, we adopt a generalized fiducial inference (GFI) approach because it has similar to an objective Bayes' interpretation with data driven priors, gives a systematic method of constructing a distribution function given a data generating equation such as a linear model, and it does not suffer from the issue of arbitrary normalizing constants which arise in many objective Bayesian priors (Berger and Pericchi [3]). In this manuscript, we will provide a gentle introduction to GFI. A fuller account of GFI is given in the recent review paper of Hannig et al. [11].

An advantage of both our approach and the nonlocal prior approach of Johnson and Rossell [13] is that in addition to providing theoretical guarantees, our statistical procedures yield estimates of the posterior distribution over subsets of covariates. This is in contrast to frequentist penalization based methods or Bayesian procedures fully dedicated to maximum a posteriori probability (MAP) estimation. Such methods do not yield the posterior probability of a chosen model (i.e., the relative probability, given the observed data, of a given model against competing models). Furthermore, Ghosh and Ghattas [10] argue that joint summaries of subsets of covariates are more robust to collinearity.

With our approach to constructing a posterior-like distribution whose probability mass function value is negligible for subsets of the parameter space which are not ε -admissible, with a typical sparsity assumption we are able to show that the probability of the true data generating model converges to 1 asymptotically in n and p . This consistency is shown to be true even with p growing almost exponentially in n . The reason being that the true model yields a stronger signal since it no longer has to compete within an overly redundant sample space.

The paper is organized as follows. Section 2 serves to introduce the general methodology and computational algorithm for carrying out our variable selection procedure based on a recent algorithm for explicit L_0 minimization (Bertsimas,

King and Mazumder [5]), which is fast enough to be used on real data. The conditions needed for the main results, and the main results are presented and discussed in Section 3. Proofs are organized in the Appendix. We demonstrate the empirical performance of our procedure and compare it to other Bayesian and frequentist methods in simulation setups on synthetic data in Section 4. Computer code implementing our procedure is provided at <https://jonathanpw.github.io/software.html>.

2. Methodology. As described in the previous section, our idea behind exploiting a nonredundancy property of the true data generating model relies on constructing a probability distribution concentrated on what we denote as ε -admissible subsets. This object is defined precisely in Definition 2.1, but first an aside on the notation used throughout the paper.

The function $|\cdot|$ denotes the absolute value function if its argument is scalar-valued, and denotes the cardinality function if its argument is set-valued. The norms $\|\cdot\|_2$ and $\|\cdot\|_0$ refer, respectfully, to the usual L_2 and L_0 norms defined on finite-dimensional Euclidean spaces. Lastly, $\|\cdot\|_F$ denotes the matrix Frobenius norm.

Let Y be an n -dimensional random vector, X an $n \times p$ matrix with columns scaled to have unit norm, and β^0 a fixed p -dimensional vector with nonzero (or active) components indexed by the subset $M_o \subset \{1, \dots, p\}$, with

$$(2) \quad Y \sim N_n(X_{M_o}\beta_{M_o}^0, (\sigma_{M_o}^0)^2 I_n).$$

The design matrix denoted by X_{M_o} is defined as the matrix composed of only those columns of X corresponding to the index set M_o . The subscript “ o ” refers to the interpretation of M_o as corresponding to the “oracle” subset of covariates. Moreover, $\beta_{M_o}^0$ denotes the true values of the oracle coefficients, while β_{M_o} is understood as a random vector whose uncertainty resides in not knowing the true coefficients $\beta_{M_o}^0$. For any subset M , the vector β_M^0 refers to the projection of the column space of X_M on the true coefficients $\beta_{M_o}^0$, that is, $\beta_M^0 = (X_M' X_M)^{-1} X_M' X_{M_o} \beta_{M_o}^0 = E_y(\widehat{\beta}_M)$, where $\widehat{\beta}_M := (X_M' X_M)^{-1} X_M' Y$. The subscript on $E_y(\cdot)$ is used to denote the expectation taken with respect to the sampling distribution of the data Y . Lastly, $\sigma_{M_o}^0 > 0$ denotes the true unknown error standard deviation, while σ_{M_o} is a random variable whose distribution expresses the uncertainty from not knowing $\sigma_{M_o}^0$, under the oracle model.

The objective is to construct a statistical procedure which can be shown, asymptotically and demonstrated empirically, to be able to identify M_o as the index set of the oracle model within the sample space of all 2^p candidate subsets $M \subset \{1, \dots, p\}$. For each index set, M , in the sample space the conditional sampling distribution of the data is assumed as

$$(3) \quad Y|\beta_M, \sigma_M^2 \sim N_n(X_M \beta_M, \sigma_M^2 I_n).$$

The centerpiece of our methodology is then the following definition.

DEFINITION 2.1. Assume fixed $\varepsilon > 0$. A given β_M coupled within some index subset $M \subset \{1, \dots, p\}$ is said to be ε -admissible if and only if $h(\beta_M) = 1$, where

$$(4) \quad h(\beta_M) := \mathbf{I}\left\{\frac{1}{2}\|X'(X_M\beta_M - Xb_{\min})\|_2^2 \geq \varepsilon\right\},$$

and b_{\min} solves

$$\min_{b \in \mathbb{R}^p} \frac{1}{2}\|X'(X_M\beta_M - Xb)\|_2^2$$

subject to

$$\|b\|_0 \leq |M| - 1.$$

Observe that this definition is consistent with the heuristic description of ε -admissible subsets given in the previous section. In particular, if the subset of covariates indexed by M is linearly dependent or if one of the components of β_M is zero, then $h(\beta_M) = 0$. The subtlety in this definition is assuming an appropriately chosen ε which is able to strike an optimal balance for distinguishing signal from noise. Intuitively, $\varepsilon = \varepsilon(n, p, M)$, that is, a function of the amount of information available given by n , the difficulty of the problem represented by p and information about a given M being considered such as $|M|$. For instance, if $|M| > n$ then $h(\beta_M) = 0$ because X_M cannot have full rank. In this case, any $\varepsilon > 0$ will work, but the choice of ε matters a lot if $|M| \leq n$. The choice of ε is a major focus of Section 3 where the main results of the paper are presented, and from where we suggest the following default choice:

$$(5) \quad \varepsilon = \Lambda_M \hat{\sigma}_M^2 \left(\frac{n^{0.51}}{9} + |M| \frac{\log(p\pi)^{1.1}}{9} - p_o \right)_+,$$

where $\Lambda_M := \|H_M X\|_F^2$ and $\hat{\sigma}_M^2 := \text{RSS}_M / (n - |M|)$ with $\text{RSS}_M := y'(I_n - H_M)y$ and $H_M := X_M(X_M'X_M)^{-1}X_M'$, and the vector y an observation from the true model (2). The parameter p_o represents prior belief about $|M_o|$, the number of covariates in the true model M_o . In practice, a value of p_o can be directly specified or selected by cross-validation. A built-in cross-validation procedure is included in the accompanying software to this paper. Details are provided with the simulation study in Section 4.

Within the h function in Definition 2.1, the quantity $\frac{1}{2}\|X'(X_M\beta_M - Xb_{\min})\|_2^2$ represents the difference in prediction for a subset M against all subsets with fewer covariates. This measure of distance has been adapted from Candes and Tao [8], but they deal with the error $\|X'(y - Xb)\|_\infty$ over $b \in \mathbb{R}^p$. This is very different from using $X_M\beta_M$ in place of y because the former results in a noiseless measure of distance. To illustrate, observe that

$$\text{E}_y(\|X'(Y - Xb)\|_2^2) = \|X'(X_{M_o}\beta_{M_o}^0 - Xb)\|_2^2 + (\sigma_{M_o}^0)^2 \cdot p.$$

There are various reasons for using the quantity $X'(X_M\beta_M - Xb)$ from the Dantzig selector (Candes and Tao [8]) versus simply the difference in predictions $(X_M\beta_M - Xb)$ as in the LASSO (Tibshirani [21]). One reason is that $X'(X_M\beta_M - Xb)$ accounts for difference in predictions as well as correlations with the explanatory data, as discussed in Berk [4]. If the difference in predictions is small but is highly correlated with the design matrix, then it is likely that the smaller subset of covariates is unable to account for the effect of one or more of the covariates in M . Thus, using $X'(X_M\beta_M - Xb)$ instead of just the difference in predictions is a method of controlling for potential omitted variable effects which could incorrectly find a close fitting subset to M . Another advantage of $X'(X_M\beta_M - Xb)$ is that it is invariant under orthogonal transformations of the design matrix, as pointed out in Candes and Tao [8].

Now that the foundation for ε -admissible subsets of the parameter space has been laid out, it remains to show how Definition 2.1 can be coupled with a likelihood based approach for constructing a probability distribution over index subsets $M \subset \{1, \dots, p\}$. This is a common strategy for Bayesian approaches, that is, construct a prior density with desirable properties for variable selection and then couple the prior with a likelihood function of the data to study the resulting posterior distribution. However, it is not clear what sort of a prior to use within our ε -admissible subsets approach, and recent developments in generalized fiducial inference (GFI) offer a systematic method of deriving objective Bayes-like posterior distributions.

To illustrate as in Hannig et al. [11], suppose that some data $Y = G(U, \theta)$ for some deterministic data generating equation $G(\cdot, \cdot)$, some parameters θ and some random component U whose distribution is independent of θ and is completely known. The generalized fiducial distribution of θ is then given by

$$r(\theta|y) = \frac{f(y, \theta)J(y, \theta)}{\int_{\Theta} f(y, \theta')J(y, \theta') d\theta'}$$

where f is the likelihood function and

$$J(y, \theta) = D\left(\frac{d}{d\theta}G(u, \theta)\Big|_{u=G^{-1}(y, \theta)}\right)$$

with $D(A) = (\det A' A)^{\frac{1}{2}}$. The component $J(y, \theta)$ is termed the *Jacobian* because it results from inverting the data generating equation on the data. We are committing a slight abuse of notation as $r(\theta|y)$ is not a conditional density in the usual sense. Instead, we are using this notation to stress that the generalized fiducial distribution is a function of the observed data y .

To make matters concrete in the linear model setting of (3), the parameters are $\theta = (\beta_M, \sigma_M)$, the data generating equation is specified as $G(U, (\beta_M, \sigma_M)) = X_M\beta_M + \sigma_M U$ where $U \sim N_n(0, I_n)$, and the Jacobian term reduces to

$J(y, (\beta_M, \sigma_M)) = \sigma_M^{-1} |\det(X'_M X_M)|^{\frac{1}{2}} \text{RSS}_M^{\frac{1}{2}}$. Thus,

$$r((\beta_M, \sigma_M)|y) \propto \sigma_M^{-n} e^{-\frac{\|y - X_M \beta_M\|_2^2}{2\sigma_M^2}} \sigma_M^{-1} |\det(X'_M X_M)|^{\frac{1}{2}} \text{RSS}_M^{\frac{1}{2}} \cdot h(\beta_M),$$

where the factor of $h(\beta_M)$ appears in the likelihood from only considering ε -admissible subsets of the parameter space. Accordingly, as is done with a Bayesian posterior density and in Section 3 of Hannig et al. [11], define the GFI probability of a given subset M to be proportional to the normalizing constant of $r((\beta_M, \sigma_M)|y)$. That is,

$$\begin{aligned} r(M|y) &:= \frac{\int f(y, (\beta_M, \sigma_M)) J(y, (\beta_M, \sigma_M)) h(\beta_M) d(\sigma_M, \beta_M)}{\sum_{j=1}^p \sum_{|M|=j} \int f(y, (\beta_M, \sigma_M)) J(y, (\beta_M, \sigma_M)) h(\beta_M) d(\sigma_M, \beta_M)} \\ &\propto \int_{\mathbb{R}^{|M|}} \int_0^\infty h(\beta_M) \frac{|\det(X'_M X_M)|^{\frac{1}{2}} \text{RSS}_M^{\frac{1}{2}}}{\sigma_M^{n+1} e^{\frac{(y - X_M \beta_M)'(y - X_M \beta_M)}{2\sigma_M^2}}} d\sigma_M d\beta_M, \end{aligned}$$

which simplifies to

$$(6) \quad r(M|y) \propto \pi^{\frac{|M|}{2}} \Gamma\left(\frac{n - |M|}{2}\right) \text{RSS}_M^{-\frac{(n - |M| - 1)}{2}} E(h(\beta_M)),$$

where the expectation is taken with respect to the location-scale multivariate T distribution,

$$(7) \quad t_{n - |M|}\left(\widehat{\beta}_M, \frac{\text{RSS}_M}{n - |M|} (X'_M X_M)^{-1}\right),$$

with $\widehat{\beta}_M := (X'_M X_M)^{-1} X'_M y$. Notice that the quantity $E(h(\beta_M))$ is a function of the observed data y .

Observe that (6) expresses the relative likelihood of the subset M over all 2^p possible subsets. The expression can be described as a product of two terms, the first being comprised of information from the sampling distribution of the data and largely driven by the residual sum of squares, RSS_M , and the second having to do with the ε -admissibility of β_M , in the form of $E(h(\beta_M))$. Thus, the support of $r(M|y)$ in (6) is dominated by the ε -admissible subsets, as desired.

Section 3 provides the conditions and supporting lemmas and theorems needed to show that $r(M_o|Y) \rightarrow 1$ in probability as $n, p \rightarrow \infty$. First however, a few remarks are provided about computing $r(M|y)$ on actual data.

2.1. *Remarks on computation.* With a probability distribution now defined over ε -admissible subsets, it must be demonstrated that $r(M|y)$ in (6) can be efficiently computed. There are two main computational issues to deal with. The first is to evaluate $h(\beta_M)$ for a given β_M , and the second is to sample subsets M via pseudo-marginal based MCMC. The computational complexity and the need for

pseudo-marginal based MCMC arises because neither $h(\beta_M)$ nor $E(h(\beta_M))$ have a closed-form solution.

To evaluate $h(\beta_M)$ for a given β_M we adapt an explicit L_0 minimization algorithm introduced in Bertsimas, King and Mazumder [5]. The authors state that their algorithm borrows ideas from projected gradient descent and methods in first-order convex optimization, and solves problems of the form $\min_{b \in \mathbb{R}^p} g(b)$ subject to $\|b\|_0 \leq \kappa$, where $g(b) \geq 0$ is convex and has Lipschitz continuous gradient: $\|\nabla g(b) - \nabla g(\tilde{b})\|_2 \leq l\|b - \tilde{b}\|_2$. The algorithm is not guaranteed to find a global optimum (unless formal optimality tests are run, which can take a long time), but Bertsimas, King and Mazumder [5] provide provable guarantees that the algorithm will converge to a first-order stationary point, which is defined as a vector $\tilde{b} \in \mathbb{R}^p$ with $\|\tilde{b}\|_0 \leq \kappa$ which satisfies $\tilde{b} = \tilde{b} - \frac{1}{l}\nabla g(\tilde{b})$. Paraphrasing from Bertsimas, King and Mazumder [5], their algorithm detects the active set after a few iterations, and then takes additional time to estimate the coefficient values to a high accuracy level. In our application of their algorithm, we are *not* first-most interested in finding a global optimum. To evaluate $h(\beta_M)$, we need only determine if $\min_{b \in \mathbb{R}^p} \frac{1}{2}\|X'(X_M\beta_M - Xb)\|_2^2$ is smaller than ε [as in (5)]. For β_M which are not ε -admissible, the objective function, $\frac{1}{2}\|X'(X_M\beta_M - Xb)\|_2^2$, can be made small very quickly via our implementation of the L_0 minimization algorithm. To illustrate how, consider the following specifics of our implementation. The precise details regarding the algorithm can be found accompanying our software documentation at <https://jonathanpw.github.io/software.html>.

First, to estimate $E(h(\beta_M))$ we use a sample mean of sample vectors drawn from the location-scale multivariate T distribution in (7). This multivariate T distribution is centered at the least squares estimator, $\hat{\beta}_M$, and multivariate theory suggests that $\hat{\beta}_M$ will on average be close to the coefficients β_M^0 . By warm starting the L_0 minimization algorithm at $\hat{\beta}_M$ with the smallest coefficient removed, subsets corresponding to β_M^0 with at least one zero coefficient typically yield $h(\beta_M) = 0$ within a few steps of the algorithm.

Second, as per the definition of $h(\cdot)$ in (4) the objective function is minimized over all $b \in \mathbb{R}^p$ with $\|b\|_0 \leq |M| - 1$. Hence, the κ required for the L_0 minimization algorithm from Bertsimas, King and Mazumder [5] is naturally chosen for us as $\kappa = |M| - 1$. Knowing how to choose κ greatly reduces the L_0 optimization problem. Moreover, our implementation is further simplified by the fact that the closest prediction to $X_M\beta_M$ for a given M is guaranteed to have $|M| - 1$ covariates. Accordingly, the objective function in $h(\beta_M)$ need *not* be minimized over all $b \in \mathbb{R}^p$ with $\|b\|_0 \leq |M| - 1$, but can be minimized over all $b \in \mathbb{R}^p$ with $\|b\|_0 = |M| - 1$.

The second computational issue is to sample subsets M via pseudo-marginal based MCMC. We do this by using the Grouped Independence Metropolis Hastings (GIMH) algorithm from Andrieu and Roberts [1], but originally introduced in Beaumont [2]. The reason standard MCMC techniques do not apply is that there is

no obvious closed-form expression for the probability mass function (6) because of the expectation, $E(h(\beta_M))$, in the expression. As described in Andrieu and Roberts [1], such situations warrant introducing a latent variable to yield analytical expressions or easier implementation.

In the case of $r(M|y)$ in (6), we introduce the latent location-scale multivariate T vector in (7) from within the expectation $E(h(\beta_M))$. Our pseudo-marginal based MCMC is carried out by sampling an index subset M along with sampling some pre-specified number, N , of multivariate T vectors (corresponding to M) from (7). The sample of multivariate T vectors, say B , is then used to compute the sample mean estimate of $E(h(\beta_M))$. Accordingly, we define a joint Markov chain on (M, B) , but discard B to obtain samples from the marginal distribution of M . As argued in Andrieu and Roberts [1], this is a valid MCMC sampling strategy, but is known to suffer from slower mixing than if we were able to integrate the β_M out of the mass function $r(M|y)$ in (6), that is, analytically evaluate $E(h(\beta_M))$. However, this is not possible given the h function in (4). Additionally, the mixing associated with pseudo-marginal approaches is known to be poor when the number of importance samples (N , the sample size of B) is small. These practical bottlenecks outline avenues for future research. Nonetheless, we demonstrate in Section 4 that our computational strategies are efficient enough to be implemented on actual data, in comparison to other common penalized likelihood and Bayesian approaches.

3. Theoretical results. The main objective of this section is to show under what conditions, asymptotically, $r(M_o|Y)$ in (6) will converge to 1, particularly if $p \gg n$. The ε -admissible subsets approach is able to achieve such a strong consistency result because the resulting sample space is effectively reduced to only those subsets with no redundancies. The essence of the mathematical result is that the space of ε -admissible sets is small enough that the true model can be detected. This addresses the issue raised in Ghosh and Ghattas [10] that high-dimensional settings often lead to arbitrarily small probabilities for all models (including the true model) simply because there are too many models to consider.

3.1. Discussion of the conditions. The first two conditions, Condition 3.1 and Condition 3.2, are to ensure that the true model, M_o , is identifiable. Observe from (4) that ε is used to control the sensitivity and specificity for identifying ε -admissible subsets. In particular, if ε is too large, then $h(\beta_{M_o})$ will incorrectly be set to zero implying that β_{M_o} is not ε -admissible. Condition 3.1 specifies how large ε can be while the true model remains identifiable. This condition turns out to be critically important in actual data applications because computing $h(\beta_{M_o})$ is closely related to the comparison in equation (8).

CONDITION 3.1. For large n and p ,

$$(8) \quad \frac{1}{18} \|X'(X_{M_o}\beta_{M_o}^0 - Xb_{\min})\|_2^2 \geq \varepsilon,$$

where b_{\min} solves

$$\min_{b \in \mathbb{R}^p} \frac{1}{2} \|X'(X_{M_o} \beta_{M_o}^0 - Xb)\|_2^2$$

subject to

$$\|b\|_0 \leq |M_o| - 1.$$

Condition 3.2 is born from Lemma A.1 which is an important necessary result for the main result of this paper, Theorem 3.9. The term $\log(n)^\gamma$ represents the sparsity assumption for the true model, that is, the number of covariates in the true model must not exceed $\log(n)^\gamma$ for some fixed scalar $\gamma > 0$. The γ parameter indicates that the asymptotic results remain true if the true model grows faster than $\log(n)$, but not faster than some power of $\log(n)$. In finite-sample applications, γ has no consequence and can be ignored.

The constant $\alpha \in (0, 1)$ reflects the only explicit restriction needed on the sample space of 2^p subsets to show that $r(M_o|Y) \rightarrow 1$ in probability for large n and p , Theorem 3.9. The residual sum of squares term in $r(M|Y)$ in (6) cannot be controlled [as a ratio to $r(M_o|Y)$] for arbitrary subsets M with $|M| = O(n)$ because the column span of X_M includes $y \in \mathbb{R}^n$ when $\text{rank}(X_M) = n$. To eliminate such subsets from the sample space, Condition 3.2 requires that only subsets of size $|M| \leq n^\alpha$ can be given nonzero probability. However, recall from Definition 2.1 that $h(\beta_M) = 0$ if $|M| > n$ because in this case the columns of X_M must be linearly dependent. Accordingly, all subsets M with $|M| > n$ are given zero probability, by definition. Evidenced by this fact, the only explicit restriction placed on the sample space is that subsets M with $|M| \in (n^\alpha, n)$ are excluded. In sparse settings, it is assumed that $|M_o| \ll n$ anyway, so neglecting such subsets is reasonable. Convergence to the true model M_o will be quicker for smaller α because there are less models to consider, but too small of an α will exclude M_o from the sample space.

In Condition 3.2, and for the remainder of this section assume that $\gamma > 0$, say $\gamma = 1$, and $\alpha \in (0, 1)$, say $\alpha = 0.5$, have been chosen and fixed at appropriate values.

CONDITION 3.2. The true model M_o satisfies $|M_o| \leq \log(n)^\gamma$, and

$$\lim_{\substack{n \rightarrow \infty \\ p \rightarrow \infty}} \min \left\{ \frac{\Delta_M}{|M_o| \log(p)} : M \neq M_o, |M| \leq |M_o| \right\} = \infty,$$

$$\liminf_{\substack{n \rightarrow \infty \\ p \rightarrow \infty}} \frac{n^{1-\alpha}}{\log(p)} > 2 \quad \text{and} \quad \log(p) < \frac{n - |M_o| - 1}{4 \log(n)^\gamma}$$

for large n and p , where $\Delta_M := \|X_{M_o} \beta_{M_o}^0 - H_M X_{M_o} \beta_{M_o}^0\|_2^2$ as in Lai, Hannig and Lee [14].

This is a slightly weaker version of condition (11) in Lai, Hannig and Lee [14]. They relate Condition 3.2 to the sparse Riesz condition (Zhang and Huang [22]) which requires that the eigenvalues of $X'_M X_M/n$ are uniformly bounded away from 0 and ∞ . Essentially, Δ_M is a measure of how distinct the true model predictions $X_{M_o} \beta_{M_o}^0$ are from their projection onto the column space of X_M for $M \neq M_o$ and $|M| \leq |M_o|$. Recall that $H_M := X_M(X'_M X_M)^{-1} X'_M$. In particular, if X_M is orthogonal to X_{M_o} , then $\Delta_M = \|X_{M_o} \beta_{M_o}^0\|_2^2$ which will be much larger than the denominator, $|M_o| \log(p)$. The requirements of this condition are reasonable because Δ_M grows very fast for M such that $M_o \not\subseteq M$.

Condition 3.2 is important for being able to identify the true model among other models M with $|M| \leq |M_o|$. The next two conditions address the requirements for M with $|M| > |M_o|$, which primarily rely on the fact that such subsets are not ε -admissible.

Conditions 3.3 and 3.4 demonstrate how large ε needs to be to achieve the consistency result of the main theorem. Condition 3.3 states that for subsets of covariates with redundancies, ε needs to be larger than the difference in projections of the true model prediction, $X_{M_o} \beta_{M_o}^0$, onto M and onto a strict subset of M . This condition facilitates the intuition that the variable selection procedure will not concentrate on subsets M with redundant covariates. If a given subset M is not ε -admissible, then the difference in projections will be small so that the condition is easily achieved.

CONDITION 3.3. For any M with $|M| > |M_o|$, for large n and p ,

$$\frac{9}{2} \|X'(H_M - H_{M(-)})X_{M_o} \beta_{M_o}^0\|_2^2 < \varepsilon,$$

where $H_{M(-)}$ is the projection matrix for M after omitting the covariate which minimizes $\|X'(H_M - H_{M(-)})X_{M_o} \beta_{M_o}^0\|_2^2$.

In fact, if $M_o \subset M$ with $|M_o| < |M|$, then $H_M X_{M_o} \beta_{M_o}^0 = H_{M(-)} X_{M_o} \beta_{M_o}^0$ in which case Condition 3.3 holds trivially.

Lastly, Condition 3.4 describes the rate at which ε needs to grow to achieve the consistency of the main result. The distinction between Condition 3.3 and Condition 3.4 is that the former provides a necessary lower bound for arguing that $E(h(\beta_M))$ vanishes for M such that $|M| > |M_o|$, while the latter provides the rate at which ε must grow to achieve the consistency result of Theorem 3.9.

CONDITION 3.4.

$$\lim_{\substack{n \rightarrow \infty \\ p \rightarrow \infty}} \min_{|M| \leq n^\alpha} \frac{\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2} + D_1 |M_o| - \varphi(M, n, p)}{\log(n)} = \infty,$$

where

$$\varphi(M, n, p) := 4e^2 n^\alpha + (D_1 + (1 + 4e^2) \log(p))|M| + \log\left(\frac{|M|}{1 - \frac{1}{n-|M|}}\right),$$

and $D_1 = \frac{1}{2} \log\left(\frac{6\pi}{1 - \frac{n^\alpha + 2}{n}}\right)$. Additionally, $\frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2} < \frac{n-|M|}{2}$ for all M with $|M| \leq n^\alpha$.

The terms which compete with ε arise in the proofs of Lemma A.1 and Theorem 3.9. Recall that $\hat{\sigma}_M^2 := \text{RSS}_M / (n - |M|)$, where RSS_M is the classical residual sum of squares for model M , and that $\Lambda_M := \|H_M X\|_F^2$. The Λ_M term arises from Lemmas 3.5 and 3.6. It is intimately related to the presence of collinearity among the covariates, and Condition 3.4 implies that ε must account for collinearity by controlling for Λ_M . Observe that if X is orthogonal, then $\Lambda_M = |M|$.

3.2. *Main result.* The first two results are lemmas which are needed in the proofs of Theorems 3.7 and 3.8. Lemma 3.5 illustrates the rate at which β_M concentrates around its mean, $\hat{\beta}_M$, the least squares estimator, and Lemma 3.6 illustrates the rate at which $\hat{\beta}_M$ concentrates around its mean, $E_y(\hat{\beta}_M)$. Theorem 3.7 uses these two lemmas to bound the rate at which β_M concentrates around $E_y(\hat{\beta}_M)$ for subsets M with $|M| > |M_o|$. This yields an upper bound on $E(h(\beta_M))$ with a probabilistic guarantee, and implies that $E(h(\beta_M))$ vanishes for large n and p , for large non- ε -admissible subsets. The proofs are relegated to the Appendix.

LEMMA 3.5. *For any fixed $c_1 \in (0, 1)$ assume $|M| \leq c_1 n$, and choose n and p such that $\frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2} < \frac{n-|M|}{2}$. If*

$$\beta_M \sim t_{n-|M|}\left(\hat{\beta}_M, \frac{\text{RSS}_M}{n - |M|} (X'_M X_M)^{-1}\right),$$

where $\hat{\beta}_M = (X'_M X_M)^{-1} X'_M y$, then

$$P\left(\frac{1}{2} \|X'_M X_M (\beta_M - \hat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) \leq \frac{2^{\frac{3}{2}} 3 |M| \hat{\sigma}_M \sqrt{\Lambda_M} e^{-\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2}}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n-|M|})}.$$

In the next lemma, P_y is used to denote the probability measure associated with the sampling distribution of the data Y .

LEMMA 3.6. *Assume $|M| < n$, and $Y | \beta_M, \sigma_M^2 \sim N_n(X_M \beta_M, \sigma_M^2 I_n)$. Then the least squares estimator $\hat{\beta}_M \sim N_{|M|}(E_y(\hat{\beta}_M), \sigma_M^2 (X'_M X_M)^{-1})$, and*

$$P_y\left(\frac{1}{2} \|X'_M X_M (\hat{\beta}_M - E_y(\hat{\beta}_M))\|_2^2 \geq \frac{\varepsilon}{9}\right) \leq \frac{3|M| \sigma_M \sqrt{\Lambda_M}}{\sqrt{\pi \varepsilon}} e^{-\frac{\varepsilon}{9\sigma_M^2 \Lambda_M}}.$$

Combining these two lemmas gives the following nonasymptotic concentration result for models that are larger than the true model. Recall that the expectation $E(h(\beta_M))$ depends on the observed data y . The following two theorems study the frequentist behavior of this quantity with respect to the sampling distribution of Y .

THEOREM 3.7. *For any fixed $c_1 \in (0, 1)$ suppose $|M_o| < |M| \leq c_1 n$, choose n and p such that $\frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2} < \frac{n-|M|}{2}$, and assume that ε satisfies Condition 3.3. Then*

$$P_y \left(E(h(\beta_M)) \leq \frac{2^{\frac{3}{2}} 3 |M| \hat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n-|M|})} e^{-\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2}} \right) \geq 1 - \frac{3 |M| \sigma_M \sqrt{\Lambda_M}}{\sqrt{\varepsilon \pi} e^{\frac{\varepsilon}{9\sigma_M^2 \Lambda_M}}}$$

The next result is a probabilistic guarantee the true model is ε -admissible given ε satisfies Condition 3.1. This result is a statement that M_o is identifiable.

THEOREM 3.8. *For any fixed $c_1 \in (0, 1)$, suppose $|M_o| \leq c_1 n$, choose n and p such that $\frac{\varepsilon}{9\Lambda_{M_o} \hat{\sigma}_{M_o}^2} < \frac{n-|M_o|}{2}$, and assume that ε satisfies Condition 3.1. Then*

$$P_y \left(E(h(\beta_{M_o})) \geq 1 - \frac{2^{\frac{3}{2}} 3 p_{M_o} \hat{\sigma}_{M_o} \sqrt{\Lambda_{M_o}}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n-p_{M_o}}) e^{\frac{\varepsilon}{18\Lambda_{M_o} \hat{\sigma}_{M_o}^2}}} \right) \geq 1 - \frac{3 p_{M_o} \sigma_{M_o} \sqrt{\Lambda_{M_o}}}{\sqrt{\varepsilon \pi} e^{\frac{\varepsilon}{9\sigma_{M_o}^2 \Lambda_{M_o}}}}$$

The following result is the main result of the paper. It shows that the ratio of the generalized fiducial probability of the true model to the sum over that of all other subsets of covariates M satisfying $|M| \leq n^\alpha$ will converge to 1 in probability for large n and p . Note that the restriction to subsets M with $|M| \leq n^\alpha$ is a stronger restriction than $|M| \leq c_1 n$, which is sufficient for Theorems 3.7 and 3.8. The reason being that the main result is stronger than the results of these two theorems. In fact, Theorems 3.7 and 3.8 are nonasymptotic results that hold for each fixed model M separately, while Theorem 3.9 is an asymptotic result which applies uniformly over all $|M| \leq n^\alpha$. Just like with a conditional distribution, $r(M|Y)$ is obtained by replacing the observed data y with the random variable Y (random with respect to the sampling distribution), in (6).

THEOREM 3.9. *Given Conditions 3.1–3.4, the true model M_o satisfies*

$$\frac{r(M_o|Y)}{\sum_{j=1}^{n^\alpha} \sum_{M:|M|=j} r(M|Y)} \xrightarrow{P_y} 1$$

as $n \rightarrow \infty$ or $n, p \rightarrow \infty$.

Although this is an asymptotic result, many of the ingredients that are used in its proof are nonasymptotic concentration results which are valid if Conditions

3.1–3.4 are satisfied. Therefore, it can be expected that when these conditions are satisfied, in finite-sample situations the generalized fiducial distribution will concentrate on the true model M_o . This expectation is indeed validated by the empirical performance of the procedure, which is demonstrated in the following section.

4. Simulation results. This section serves to demonstrate the empirical performance of our algorithm on synthetic data. It is comprised of essentially two simulation setups. The first setup, similar to that presented in Johnson and Rossell [13], compares our procedure to the nonlocal prior approach, the spike and slab LASSO of Ročková and George [18], the elastic net as implemented in the *Python* module `scikit-learn` (Pedregosa et al. [17]), and to the SCAD as implemented in the *R* package `ncvreg` (Breheny and Huang [7]). The authors of the nonlocal prior and the spike and slab LASSO, respectively, have made available the *R* packages `mombf` and `SSL` for implementing their methods.

The second setup illustrates a critical difference between our ε -admissible subsets and the nonlocal prior approach. Namely, for highly collinear finite-sample settings in which the true model is not uniquely expressed, given the level of noise in the data (i.e., $\sigma_{M_o}^0$), we demonstrate that our approach concentrates (in the sense of the MAP estimator) on subsets with fewer covariates without sacrificing prediction error. The intuition for why this should be the case was discussed in Section 1.

4.1. *Simulation setup 1.* Here, we generate 2000 data vectors y according to model (2) with M_o consisting of 8 covariates corresponding to $\beta_{M_o}^0 = (-1.5, -1, -0.8, -0.6, 0.6, 0.8, 1, 1.5)'$, and $\sigma_{M_o}^0 = 1$. The $n \times p$ design matrix X is generated with rows from the $N_p(0, \Sigma)$ distribution, where the diagonal components $\Sigma_{ii} = 1$ and the off-diagonal components $\Sigma_{ij} = \rho$ for $i \neq j$. The first 1000 y correspond to an independent design with $\rho = 0$, while the last 1000 y correspond to $\rho = 0.25$, as in the simulation setup of Johnson and Rossell [13]. Note that 2000 design matrices X are generated, and one y is generated from each design. The sample size n is set at $n = 100$, and $p = 100, 200, 300, 400, 500$ are all considered.

We implement our algorithm on each of the 2000 synthetic data sets for 15,000 MCMC steps with the first 5000 discarded. Squared coefficient estimates from elastic net (using `scikit-learn`) added by n^{-2} serve as MCMC covariate proposal weights. The default ε in (5) is used, and we implemented a 10-fold cross-validation scheme for choosing our tuning parameter p_o .

The cross-validation consists of breaking the data into 10 folds (with a different set of 10 observations held out at each fold since $n = 100$), and implementing our MCMC algorithm separately for each p_o in the grid $\{1, 2, \dots, 10\}$, on each of the 10 training sets. Each of the 10 implementations of the MCMC on each of the 10 training sets is run for 200 steps with the first 100 steps discarded ($N = 30$ is set during the cross-validation procedure). Squared nonzero coefficient estimates from

elastic net (using `scikit-learn`) serve as MCMC covariate proposal weights within the cross-validation procedure. The MAP estimated subset for each MCMC chain is used to compute the Bayesian information criterion (BIC) on the held-out test set, and the computed BIC values are averaged over the 10 test sets, for each $p_o \in \{1, 2, \dots, 10\}$. The p_o corresponding to the minimum average test set BIC is then selected.

Finally, for our implementation of the algorithm post-selection of p_o , the number of importance samples for estimating $E(h(\beta_M))$ within each step of the algorithm is set at $N = 100$ which, through empirical experimentation, seems to be enough. All competing variable selection procedures are implemented using existing software at default specifications. The one exception is that the nonlocal prior procedure is set to run for 5000 steps, as is the case in the simulation setup of Johnson and Rossell [13]. The nonlocal prior procedure/software did not scale well for increased p , and required over a weeks worth of parallel computations on a computing cluster to obtain the results for the first simulation setup. The tuning parameters for all methods are chosen with the default cross-validation procedures provided with the software. Lastly, as in the simulation section for Ročková and George [18] their λ_1 is set at 1 (with a grid of 20 λ_0 values ending at 50), and their adaptive (best performing) procedure is used with $\theta \sim \text{Beta}(1, p)$.

Figure 1 shows results of the first simulation setup. The first row of plots displays the average generalized fiducial probability of the true model (i.e., average $r(M_o|y)$), or the average posterior probability of the true model for the Bayesian nonlocal procedure (i.e., average $P(M_o|y)$), over the 1000 synthetic data sets (for $\rho = 0$ and $\rho = 0.25$, resp.). Conditional on the data, these plots address the consistency of the procedures with respect to the uncertainty from not knowing M_o . This generalized fiducial or Bayesian-like consistency is that which is dealt with in Theorem 3.9. Note that frequentist and MAP estimators do not yield posterior probability estimates, and thus cannot be compared to in the first row of plots.

The second row of Figure 1 shows the average proportion of correct model selections over the 1000 synthetic data sets (for $\rho = 0$ and $\rho = 0.25$, resp.). For the GFI and the Bayesian procedures, the MAP subset is taken to be the estimator of the true model, and in the frequentist procedures the estimated model is considered to be the subset of covariates with nonzero coefficient estimates. These plots address the consistency of the procedures with respect to repeated sampling (i.e., frequentist) uncertainty. Finally, the third row of Figure 1 presents the average root mean squared error (RMSE) over the 1000 synthetic data sets (for $\rho = 0$ and $\rho = 0.25$, resp.). The MAP estimated model is used to compute the RMSE for the GFI and Bayesian procedures. For all procedures, the RMSE is computed on an out-of-sample test set of 100 observations.

Note that the criteria used in the first two rows of Figure 1 are very strict. They only reflect instances when the procedures are exactly correct, and count the procedure as incorrect if it is missing even one covariate from the true model or includes even one spurious covariate. Often the elastic net and SCAD are able to

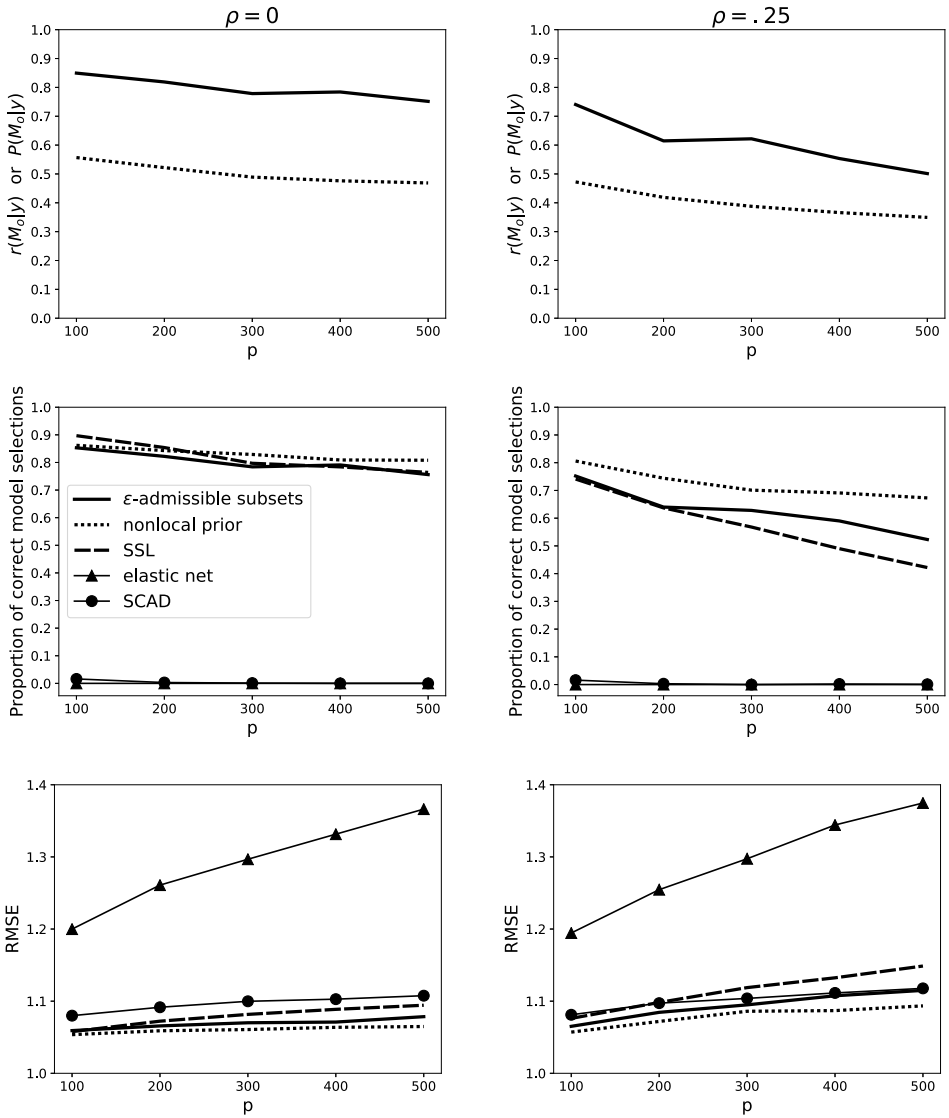


FIG. 1. The average $r(M_0|y)$, or average $P(M_0|y)$ is displayed in the first row, the average proportion of correct model selections in the second row and the average RMSE in the third. Averages are over 1000 synthetic data sets (for $\rho = 0$ and $\rho = 0.25$, resp.). For the GFI and Bayesian procedures, the MAP subset is used as the estimator of the true model, and in the frequentist procedures the estimated model is considered to be the subset of covariates with nonzero coefficient estimates. The RMSE is computed on an out-of-sample test set of 100 observations.

identify all of the true covariates but estimate extra coefficients to be nonzero. This results in poor identification of the true subset, and worse out-of-sample prediction error. The remaining procedures (including our ε -admissible subsets method) struggle to identify the two covariates with smallest coefficient magnitudes, but typically do not introduce more than one or two false positives.

Our ε -admissible subsets procedure evidently performs the best at assigning highest posterior probability to the true model. And even in comparison to the frequentist-oriented metrics, proportion of correct model selections and out-of-sample RMSE, our ε -admissible subsets procedure performs more or less on par with the best performing methods considered.

In collinear design settings (i.e., $\rho > 0$), it may be the case that a strict subset of the true data generating model is identified as the “true” model within the ε -admissible framework. This is meaningful because it manifests the fact that the true model may not be minimal (i.e., contains redundant predictors) in collinear, finite-sample settings. Furthermore, it explains the larger difference in proportion of true model selections between the ε -admissible and nonlocal prior performances in the $\rho = 0.25$ case (versus the $\rho = 0$ case), which is accompanied by only a very small change in the difference between RMSE performance. An interpretation is that for highly correlated covariates the ε -admissible subsets method may not recover the correct set of predictors (or does so less frequently), but instead it recovers a smaller set that results in almost the same predictions (as can be seen from the RMSE in Figure 1). This phenomenon is illustrated in a more extreme case of collinearity in the next simulation setup.

4.2. *Simulation setup 2.* The ε -admissible subsets approach has been developed in this paper as a method of obtaining a posterior-like distribution which effectively eliminates (i.e., assigns negligible probability to) all subsets with redundancies. To illustrate that our constructed methods do just that, consider the following setup in which the true data generating model lacks uniqueness for the small sample size $n = 30$:

$$(9) \quad Y \sim N_n(1 \cdot x^{(1)} + 1 \cdot x^{(2)} + \dots + 1 \cdot x^{(9)}, I_n),$$

where $x^{(1)}, x^{(2)}, x^{(3)} \stackrel{\text{i.i.d.}}{\sim} N_n(0, I_n)$, and

$$\begin{aligned} x^{(4)} &\sim N_n(0.25 \cdot x^{(1)}, 0.1^2 I_n), \\ x^{(5)} &\sim N_n(0.5 \cdot x^{(2)}, 0.1^2 I_n), \\ x^{(6)} &\sim N_n(-0.75 \cdot x^{(3)}, 0.1^2 I_n), \\ x^{(7)} &\sim N_n(x^{(1)} + x^{(3)}, 0.1^2 I_n), \\ x^{(8)} &\sim N_n(x^{(2)} - x^{(3)}, 0.1^2 I_n), \\ x^{(9)} &\sim N_n(x^{(1)} + x^{(2)} + x^{(3)}, 0.1^2 I_n). \end{aligned}$$

TABLE 1

The average number of covariates in the MAP estimator, M_{MAP} , ($|M_{\text{MAP}}|$) is presented in the first column, the average RMSE in the second and the average $r(M_{\text{MAP}}|y)$ or $P(M_{\text{MAP}}|y)$ in the third. Averages are over 1000 synthetic data sets from model (9). The RMSE is computed on an out-of-sample test set of 30 observations

	Model size	RMSE	$r(M_{\text{MAP}} y)$ or $P(M_{\text{MAP}} y)$
ε -admissible subsets	3.476	1.138	0.365
Nonlocal prior	8.997	1.197	0.333

With standard deviations of 0.1 and a model error standard deviation of 1, covariates $x^{(4)}, \dots, x^{(9)}$ can all be approximately expressed as a linear combination of $x^{(1)}, x^{(2)}, x^{(3)}$. Accordingly, with a small increase in error variance, model (9) can be reexpressed using various combinations of the 9 predictors. However, observe that a subset with 4 or more predictors predominantly contains redundant information.

Recall from Section 1 that the nonlocal prior approach of Johnson and Rossell [13] is designed to assign negligible probabilities to subsets containing predictor(s) with coefficients of zero. So, in theory, the full subset $\{x^{(1)}, \dots, x^{(9)}\}$ will remain the best candidate for the true model within the nonlocal prior framework. In fact, this is demonstrated to be the case in Table 1 which shows the performance of both the nonlocal prior and the ε -admissible subsets approach on 1000 data vectors y generated according to (9), with each covariate having a “true” coefficient of 1. Note that as in the first simulation setup 1000 design matrices X are generated to generate the 1000 y vectors.

Table 1 shows that the MAP estimate for the ε -admissible subsets approach contains 3–4 covariates, on average, and that in fact the average RMSE is smaller than that of the nonlocal prior approach. Indeed, the MAP estimates for the nonlocal prior procedure typically includes all 9 covariates even though the y vectors can be mostly explained by only 3 of the predictors. This simple simulation illustrates a pivotal difference between the nonlocal prior and ε -admissible subsets approaches. With $p = 9$, the implication of not discriminating against redundant subsets may seem trivial. However, the 2^p size of the sample space grows rapidly in p , and thus, puts exponentially more burden on procedures which do not discriminate based on redundancies.

5. Concluding remarks. In this paper, we have developed a new perspective for variable selection to exploit a nonredundancy property of a true data generating model. The basic idea calls for defining a true model as one which contains minimal amount of information necessary for explaining and/or predicting the observed data. The difference between our definition of a true model and the usual definition arises only in finite-sample applications, and was illustrated in Section 4.2. Within

our variable selection framework, this definition allows us to show that under a typical sparsity assumption the posterior-like probability of the true model converges to 1 asymptotically, even with p growing almost exponentially in n , with the intuition that redundancies in the sample space are very effectively eliminated. Moreover, our empirical simulation results are consistent with this strong consistency result, and as desired, it is demonstrated in a situation of high collinearity that the ε -admissible subsets approach yields a posterior-like distribution which is concentrated over subsets with fewer covariates, without sacrificing prediction error.

A nonredundancy property of a true data generating model is seemingly general enough to extend to variable or feature selection problems beyond the linear model setting, but would seem infeasible if it could not be developed for the linear model setting. Thus, the goal of this paper has been to establish the potential feasibility of exploiting such a property. In future work, we hope to extend our methods to more complicated settings.

APPENDIX: PROOFS

PROOF OF LEMMA 3.5. From the distributional assumption on β_M , it follows that

$$T := \sqrt{\frac{n - |M|}{\text{RSS}_M}} (X'_M X_M)^{\frac{1}{2}} (\beta_M - \hat{\beta}_M) \sim t_{n-|M|}(0, I_{|M|}).$$

Thus,

$$\|X'X_M(\beta_M - \hat{\beta}_M)\|_2^2 = \|AT\|_2^2 \frac{\text{RSS}_M}{n - |M|} = \|QDW'T\|_2^2 \frac{\text{RSS}_M}{n - |M|},$$

where $A = X'X_M(X'_M X_M)^{-\frac{1}{2}}$ is a $p \times |M|$ matrix which has the singular value decomposition $A = QDW'$ for Q and W each orthogonal matrices. By the definition of the multivariate T distribution,

$$T = \sqrt{\frac{n - |M|}{V}} Z \implies W'T = \sqrt{\frac{n - |M|}{V}} W'Z =: \sqrt{\frac{n - |M|}{V}} \tilde{Z} =: \tilde{T},$$

where $V \sim \chi_{n-|M|}^2$ and $Z, \tilde{Z} \sim N_{|M|}(0, I_{|M|})$. Then

$$\|AT\|_2^2 = \|QD\tilde{T}\|_2^2 = \tilde{T}'D'D\tilde{T} = \sum_1^{|M|} \tilde{T}_i^2 \lambda_i \leq \Lambda_M \max \tilde{T}_j^2,$$

where $\Lambda_M = \sum_1^{|M|} \lambda_i$, and λ_i is the i th eigenvalue of $A'A$. Observe that Λ_M also has the more intuitive expression $\Lambda_M = \text{tr}(A'A) = \|H_M X\|_F^2$.

Recall that $\hat{\sigma}_M^2 := \text{RSS}_M / (n - |M|)$. Then

$$\begin{aligned} P\left(\frac{1}{2}\|X'X_M(\beta_M - \hat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) &= P\left(\frac{1}{2}\|AT\|_2^2 \frac{\text{RSS}_M}{n - |M|} \geq \frac{\varepsilon}{9}\right) \\ &\leq P\left(\frac{1}{2}\Lambda_M \max \tilde{T}_j^2 \frac{\text{RSS}_M}{n - |M|} \geq \frac{\varepsilon}{9}\right) \\ &= P\left(\frac{n - |M|}{V} \max Z_j^2 \hat{\sigma}_M^2 \geq \frac{2\varepsilon}{9\Lambda_M}\right). \end{aligned}$$

Since $V \sim \chi_{n-|M|}^2$,

$$\begin{aligned} &P\left(\frac{1}{2}\|X'X_M(\beta_M - \hat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) \\ &\leq \int_0^\infty P\left(\max |Z_j| \geq \frac{\sqrt{2\varepsilon v / (n - |M|)}}{3\hat{\sigma}_M \sqrt{\Lambda_M}}\right) \frac{v^{\frac{n-|M|}{2}-1} e^{-\frac{v}{2}}}{2^{\frac{n-|M|}{2}} \Gamma(\frac{n-|M|}{2})} dv \\ (10) \quad &\leq \int_0^\infty \sum_{j=1}^{|M|} P\left(|Z_j| \geq \frac{\sqrt{2\varepsilon v / (n - |M|)}}{3\hat{\sigma}_M \sqrt{\Lambda_M}}\right) \frac{v^{\frac{n-|M|}{2}-1} e^{-\frac{v}{2}}}{2^{\frac{n-|M|}{2}} \Gamma(\frac{n-|M|}{2})} dv \\ &\leq \frac{3|M|\hat{\sigma}_M \sqrt{\Lambda_M(n - |M|)}}{\sqrt{2\pi\varepsilon} \left(1 + \frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2 \left(\frac{n-|M|}{2}\right)}\right)^{\frac{n-|M|}{2} - \frac{3}{2}}} \cdot \frac{\Gamma(\frac{n-|M|-1}{2})}{\Gamma(\frac{n-|M|}{2})}, \end{aligned}$$

where the last inequality follows because for the standard normal CDF, Φ for $x > 0$, $\Phi(-x) \leq \frac{1}{x\sqrt{2\pi}} e^{-\frac{x^2}{2}}$.

To simplify the bound, observe first that Jameson [12] gives

$$(11) \quad \frac{\Gamma(\frac{n-|M|-1}{2})}{\Gamma(\frac{n-|M|}{2})} \leq \frac{\sqrt{2(n - |M|)}}{n - |M| - 1}.$$

Second, observe that for $0 \leq x \leq n$,

$$\left(1 + \frac{x}{n}\right)^{-n} \leq e^{-x + \frac{x^2}{2n}} = e^{-x(1 - \frac{x}{2n})} \leq e^{-\frac{x}{2}}.$$

By assumption $|M| \leq c_1 n$, and $\frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2} < \frac{n-|M|}{2}$ which implies that

$$(12) \quad \left(1 + \frac{\varepsilon}{9\Lambda_M \hat{\sigma}_M^2 \left(\frac{n-|M|}{2}\right)}\right)^{-\frac{n-|M|}{2} + \frac{3}{2}} \leq 2^{\frac{3}{2}} e^{-\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2}}.$$

Therefore, applying (11) and (12) to (10) gives

$$P\left(\frac{1}{2}\|X'X_M(\beta_M - \hat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) \leq \frac{2^{\frac{3}{2}} 3|M|\hat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi\varepsilon} \left(1 - \frac{1}{n-|M|}\right)} e^{-\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2}}. \quad \square$$

PROOF OF LEMMA 3.6. Similar to the proof of Lemma 3.5. \square

PROOF OF THEOREM 3.7. Recall that

$$h(\beta_M) = \mathbf{I}\left\{\frac{1}{2}\|X'(X_M\beta_M - Xb_{\min})\|_2^2 \geq \varepsilon\right\},$$

where b_{\min} solves $\min_{b \in \mathbb{R}^p} \frac{1}{2}\|X'(X_M\beta_M - Xb)\|_2^2$ subject to $\|b\|_0 \leq |M| - 1$. Observe that

$$\begin{aligned} \mathbf{E}(h(\beta_M)) &= P\left(\frac{1}{2}\|X'(X_M\beta_M - Xb_{\min})\|_2^2 \geq \varepsilon\right) \\ (13) \quad &\leq P\left(\frac{1}{2}\|X'(X_M\beta_M - X\mathbf{E}_y(\widehat{\beta}_{M(-1)}))\|_2^2 \geq \varepsilon\right), \end{aligned}$$

where $\widehat{\beta}_{M(-1)}$ is the least squares estimate corresponding to the subset of covariates M with one covariate removed so that $\|\mathbf{E}_y(\widehat{\beta}_{M(-1)})\|_0 \leq |M| - 1$. The covariate removed is chosen to correspond to the smallest (in magnitude) component of $\widehat{\beta}_M$. To refine the bound on $\mathbf{E}(h(\beta_M))$, decompose the last expression in (13) using the triangle inequality as follows:

$$\begin{aligned} &P\left(\frac{1}{2}\|X'(X_M\beta_M - X_{M(-1)}\mathbf{E}_y(\widehat{\beta}_{M(-1)}))\|_2^2 \geq \varepsilon\right) \\ (14) \quad &\leq P\left(\frac{1}{2}\|X'X_M(\beta_M - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) \\ &\quad + \mathbf{I}\left\{\frac{1}{2}\|X'X_M(\widehat{\beta}_M - \mathbf{E}_y(\widehat{\beta}_M))\|_2^2 \geq \frac{\varepsilon}{9}\right\} \\ &\quad + \mathbf{I}\left\{\frac{1}{2}\|X'(H_M - H_{M(-1)})X_{M_o}\beta_{M_o}^0\|_2^2 \geq \frac{\varepsilon}{9}\right\}, \end{aligned}$$

where $H_M := X_M(X'_M X_M)^{-1} X'_M$. Observe that by Lemma 3.5 the first term on the right-hand side,

$$P\left(\frac{1}{2}\|X'X_M(\beta_M - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\right) \leq \frac{2^{\frac{3}{2}}3|M|\widehat{\sigma}_M\sqrt{\Lambda_M}}{\sqrt{\pi\varepsilon}(1 - \frac{1}{n-|M|})} e^{-\frac{\varepsilon}{18\Lambda_M\widehat{\sigma}_M^2}}.$$

Note that the last two terms are written as an indicator function because the uncertainty here comes from β_M . However, the second term results from the uncertainty in observing Y . Accordingly, by Lemma 3.6,

$$P_y\left(\mathbf{I}\left\{\frac{1}{2}\|X'X_M(\widehat{\beta}_M - \mathbf{E}_y(\widehat{\beta}_M))\|_2^2 \geq \frac{\varepsilon}{9}\right\} = 0\right) \geq 1 - \frac{3|M|\sigma_M\sqrt{\Lambda_M}}{\sqrt{\varepsilon\pi}e^{\frac{\varepsilon}{9\sigma_M^2\Lambda_M}}}.$$

Recall that P_y is used to denote the probability measure associated with the sampling distribution of Y . The third term in (14) must be zero by Condition 3.3. \square

PROOF OF THEOREM 3.8. Recall that

$$h(\beta_{M_o}) = \mathbb{I} \left\{ \frac{1}{2} \|X'(X_{M_o}\beta_{M_o} - X\tilde{b}_{\min})\|_2^2 \geq \varepsilon \right\},$$

where \tilde{b}_{\min} solves $\min_{b \in \mathbb{R}^p} \frac{1}{2} \|X'(X_{M_o}\beta_{M_o} - Xb)\|_2^2$ subject to $\|b\|_0 \leq |M_o| - 1$. To show the desired result, let b_{\min} be the solution to $\min_{b \in \mathbb{R}^p} \|X'(X_{M_o}\beta_{M_o}^0 - Xb)\|_2^2$ subject to $\|b\|_0 \leq |M_o| - 1$. Then observe that

$$\begin{aligned} \|X'(X_{M_o}\beta_{M_o}^0 - Xb_{\min})\|_2 &\leq \|X'(X_{M_o}\beta_{M_o}^0 - X\tilde{b}_{\min})\|_2 \\ &\leq \|X'X_{M_o}(\beta_{M_o}^0 - \beta_{M_o})\|_2 \\ &\quad + \|X'(X_{M_o}\beta_{M_o} - X\tilde{b}_{\min})\|_2. \end{aligned}$$

Note the difference between b_{\min} and \tilde{b}_{\min} here. The rightmost term is the quantity of interest because it will become $E(h(\beta_{M_o}))$ in the next few steps. The term on the left of the inequality corresponds to the quantity in Condition 3.1.

Adding and subtracting $\hat{\beta}_{M_o}$ inside the first term on the right-hand side of the second inequality, and applying the triangle inequality gives

$$\begin{aligned} &\mathbb{I} \left\{ \frac{1}{2} \|X'(X_{M_o}\beta_{M_o}^0 - Xb_{\min})\|_2^2 \geq 9\varepsilon \right\} \\ &\leq P \left(\frac{1}{2} \|X'X_{M_o}(\beta_{M_o} - \hat{\beta}_{M_o})\|_2^2 \geq \varepsilon \right) \\ &\quad + \mathbb{I} \left\{ \frac{1}{2} \|X'X_{M_o}(\hat{\beta}_{M_o} - E_y(\hat{\beta}_{M_o}))\|_2^2 \geq \varepsilon \right\} + E(h(\beta_{M_o})), \end{aligned}$$

and by applying Lemma 3.5,

$$\begin{aligned} &\mathbb{I} \left\{ \frac{1}{2} \|X'(X_{M_o}\beta_{M_o}^0 - Xb_{\min})\|_2^2 \geq 9\varepsilon \right\} \\ &\leq \frac{2^{\frac{3}{2}} 3 p_{M_o} \hat{\sigma}_{M_o} \sqrt{\Lambda_{M_o}}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n - p_{M_o}})} e^{-\frac{\varepsilon}{18 \Lambda_{M_o} \hat{\sigma}_{M_o}^2}} \\ &\quad + \mathbb{I} \left\{ \frac{1}{2} \|X'X_{M_o}(\hat{\beta}_{M_o} - E_y(\hat{\beta}_{M_o}))\|_2^2 \geq \frac{\varepsilon}{9} \right\} + E(h(\beta_{M_o})), \end{aligned}$$

where the middle term is written as an indicator function because the uncertainty here comes from β_{M_o} . This indicator is 0 by Lemma 3.6 with probability exceeding (15). Therefore, since Condition 3.1 implies that the indicator on the left side of the inequality is 1,

$$1 - \frac{2^{\frac{3}{2}} 3 p_{M_o} \hat{\sigma}_{M_o} \sqrt{\Lambda_{M_o}}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n - p_{M_o}})} e^{-\frac{\varepsilon}{18 \Lambda_{M_o} \hat{\sigma}_{M_o}^2}} \leq E(h(\beta_{M_o}))$$

with probability exceeding

$$(15) \quad 1 - \frac{3p_{M_o} \sigma_{M_o} \sqrt{\Lambda_{M_o}}}{\sqrt{\varepsilon \pi}} e^{-\frac{\varepsilon}{9\sigma_{M_o}^2 \Lambda_{M_o}}},$$

due to the uncertainty in observing Y . \square

The following lemma is needed for the proof of the main result, Theorem 3.9.

LEMMA A.1. *Assume all conditions and notation of Theorem 3.9, and without loss of generality suppose $\sigma_{M_o}^0 = 1$, where $\sigma_{M_o}^0$ is the true but unknown error standard deviation. Then the following holds.*

Case 1. *This case pertains to subsets M with $|M| \leq |M_o|$.*

$$P_y \left(\bigcap_{j=1}^{|M_o|} \bigcap_{\mathcal{M}_j} \left\{ Y : \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} \leq e^{-2 \log(n)^\gamma \log(p)} \right\} \right) \geq 1 - V_1,$$

where $\mathcal{M}_j := \{M \neq M_o : |M| = j\}$,

$$V_1 := \max_{\substack{M \neq M_o \\ |M| \leq |M_o|}} \left\{ \frac{12|M_o| e^{-\frac{\Delta_M}{72} + |M_o| \log(p)}}{\sqrt{2\pi \Delta_M}} + |M_o| e^{-\frac{\Delta_M}{12} + \frac{|M_o|}{2} + |M_o| \log(p)} \right. \\ \left. + |M_o| e^{-\frac{\xi_{n,|M_o|}}{48} \frac{(n-|M_o|-1)\Delta_M + n-|M_o|}{\log(n)^\gamma \log(p)} + |M_o| \log(p)} \right\},$$

and $\xi_{n,j} = 1 - \frac{2 \log(n)^\gamma \log(p)}{(n-j-1)/2} \rightarrow 1$.

Case 2. *This case pertains to subsets M with $|M_o| + 1 \leq |M| \leq n^\alpha$.*

$$P_y \left(\bigcap_{j=|M_o|+1}^{n^\alpha} \bigcap_{M:|M|=j} \left\{ Y : \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} \leq e^{e^2(n^\alpha + j \log(p))} \right\} \right) \geq 1 - V_2,$$

where

$$V_2 := \frac{n^\alpha}{e^{b_n [2n^\alpha - \frac{|M_o|}{2b_n} + (2 - \frac{1}{b_n} - \frac{1}{2b_n \log(p)}) (|M_o| + 1) \log(p)]}} + \frac{n^\alpha e^{-\frac{n-n^\alpha-|M_o|}{2} + n^\alpha \log(p)}}{\sqrt{\pi (n - n^\alpha - |M_o|)}},$$

and $b_n := \frac{n-n^\alpha-|M_o|}{n-|M_o|-2} \rightarrow 1$.

Note that V_1 and V_2 both vanish for large n and p by Condition 3.2. This condition also ensures that $\xi_{n,j} \in (0, 1)$ which is needed in the proof of this lemma.

PROOF OF LEMMA A.1. There are two cases to consider.

Case 1. For $M \in \mathcal{M}_j$ with $j \leq |M_o|$, let $\xi_{n,j} = 1 - \frac{2 \log(n)^\gamma \log(p)}{(n-j-1)/2} \in (0, 1)$ by Condition 3.2 for large n and p . Then

$$(16) \quad P_y \left(\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}} \right) \\ = P_y(U'(I_n - H_{M_o})U/\xi_{n,j} > \Delta_M + 2\sqrt{\Delta_M}Z + U'(I_n - H_M)U)$$

since by assumption $Y = X_{M_o}\beta_{M_o}^0 + U$ with $U \sim N_n(0, I_n)$, and so

$$\text{RSS}_{M_o} = U'(I_n - H_{M_o})U \quad \text{and} \\ \text{RSS}_M = \Delta_M + 2\beta_{M_o}^{0'} X'_{M_o} (I_n - H_M)U + U'(I_n - H_M)U \\ = Z_{\Delta_M} + U'(I_n - H_M)U,$$

where $Z_{\Delta_M} := \Delta_M + 2\sqrt{\Delta_M}Z$, and $Z \sim N(0, 1)$. Recall that $\Delta_M := \beta_{M_o}^{0'} X'_{M_o} (I_n - H_M) X_{M_o} \beta_{M_o}^0$. Continuing in (16) by subtracting $\chi_{n-|M_o|}^2$ from both sides of the inequality gives

$$(17) \quad P_y \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > \xi_{n,j} \right) = P_y \left(\frac{\chi_{n-|M_o|}^2}{\xi_{n,j}} - \chi_{n-|M_o|}^2 > Z_{\Delta_M} + \chi_{n-j}^2 - \chi_{n-|M_o|}^2 \right) \\ = P_y(\chi_{n-|M_o|}^2(1/\xi_{n,j} - 1) > Z_{\Delta_M} + \chi_{|M_o|}^2 - \chi_j^2) \\ \leq P_y(\chi_{n-|M_o|}^2(1/\xi_{n,j} - 1) > Z_{\Delta_M} - \chi_j^2).$$

The last inequality follows because the $\chi_{|M_o|}^2$ random variable is nonnegative, and removing it simplifies the remaining argument. Then

$$P_y \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > \xi_{n,j} \right) \leq P_y(\xi_{n,j} \chi_j^2 + \chi_{n-|M_o|}^2(1 - \xi_{n,j}) - 2\xi_{n,j} \sqrt{\Delta_M}Z > \xi_{n,j} \Delta_M) \\ \leq P_y(|Z| > \sqrt{\Delta_M}/6) + P_y(\chi_j^2 > \Delta_M/3) \\ + P_y \left(\chi_{n-|M_o|}^2 > \frac{\xi_{n,j} \Delta_M}{3(1 - \xi_{n,j})} \right).$$

For the second and third term, apply the Chernoff bound, and evaluate the moment generating function for the χ_j^2 and $\chi_{n-|M_o|}^2$ distributions at $1/4$. Accordingly,

$$P_y \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > \xi_{n,j} \right) \leq 2P_y(Z < -\sqrt{\Delta_M}/6) + e^{-\frac{\Delta_M}{12} + \frac{j}{2}} + e^{-\frac{\xi_{n,j} \Delta_M}{12(1-\xi_{n,j})} + \frac{n-|M_o|}{2}}.$$

Finally, the remaining probability can be controlled by the bound for the CDF of a standard normal random variable for $x > 0$, $\Phi(-x) \leq \frac{1}{x\sqrt{2\pi}} e^{-\frac{x^2}{2}}$. Hence,

$$P_y \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > \xi_{n,j} \right) \leq \frac{12e^{-\frac{\Delta_M}{72}}}{\sqrt{2\pi} \Delta_M} + e^{-\frac{\Delta_M}{12} + \frac{|M_o|}{2}} + e^{-\frac{\xi_{n,j} |M_o|}{48} \frac{(n-|M_o|-1)\Delta_M}{\log(n)^\gamma \log(p)} + \frac{n-|M_o|}{2}},$$

where the last inequality follows by observing that $j \leq |M_o|$, and recalling the expression for $\xi_{n,j}$.

Therefore, the probability that $\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}$ is satisfied for some M with $|M| \leq |M_o|$ is

$$\begin{aligned} P_y \left(\bigcup_{j=1}^{|M_o|} \bigcup_{\mathcal{M}_j} \left\{ \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}} \right\} \right) \\ \leq \sum_{j=1}^{|M_o|} \binom{p}{j} \max_{\mathcal{M}_j} P_y \left(\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}} \right). \end{aligned}$$

Note that

$$(18) \quad \binom{p}{j} = \frac{p(p-1)\cdots(p-j+1)}{j!} = \frac{p^j (1 - \frac{1}{p}) \cdots (1 - \frac{j-1}{p})}{j!} \leq p^j.$$

In fact, Luo and Chen [15] show that if $\log(j)/\log(p) \rightarrow \delta$ as $p \rightarrow \infty$, for some $\delta > 0$, then $\log \binom{p}{j} = j \log(p)(1 - \delta)(1 + o(1))$. Thus,

$$\begin{aligned} P_y \left(\bigcup_{j=1}^{|M_o|} \bigcup_{\mathcal{M}_j} \left\{ \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}} \right\} \right) \\ \leq \sum_{j=1}^{|M_o|} \max_{\mathcal{M}_j} \left\{ \frac{12e^{-\frac{\Delta_M}{72} + j \log(p)}}{\sqrt{2\pi \Delta_M}} + e^{-\frac{\Delta_M}{12} + \frac{|M_o|}{2} + j \log(p)} \right. \\ \left. + e^{-\frac{\xi_{n,|M_o|} (n-|M_o|-1)\Delta_M}{48 \log(n)^\gamma \log(p)} + \frac{n-|M_o|}{2} + j \log(p)} \right\}. \end{aligned}$$

Since $\xi_{n,|M_o|} \rightarrow 1$, this bound vanishes by Condition 3.2. Therefore, with probability exceeding one minus the above bound,

$$\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} \leq \left(1 - \frac{2 \log(n)^\gamma \log(p)}{(n-j-1)/2} \right)^{\frac{n-j-1}{2}} \leq e^{-2 \log(n)^\gamma \log(p)}$$

uniformly over all M such that $|M| \leq |M_o|$.

Case 2. Consider any subset M with $|M_o| < |M| \leq n^\alpha$ for some positive constant $\alpha < 1$, and let $\{a_n\}$ be an arbitrarily sequence of numbers. To begin, repeating the steps in (17), but subtracting $\chi_{n-j}^2/\xi_{n,j}$ on both sides instead of $\chi_{n-|M_o|}^2$, and replacing the label $\xi_{n,j}$ with a_n , yields

$$P_y \left(\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} \right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}} \right) \leq P_y(\chi_j^2 > a_n Z_{\Delta_M} + \chi_{n-j}^2(a_n - 1)),$$

where $Z_{\Delta_M} = \Delta_M + 2\sqrt{\Delta_M}Z$, $Z \sim N(0, 1)$, and $\Delta_M = \beta_{M_o}^{0'} X'_{M_o} (I_n - H_M) \times X_{M_o} \beta_{M_o}^0$. Since $M_o \subset M$ implies $\Delta_M = 0$, the above bound can be simplified by

including in the subset M any covariates in M_o not already included in M . Accordingly, let $M' := M \cup M_o$ which includes $j+l$ covariates, where $l \in \{0, \dots, |M_o|\}$ is the number of covariates not shared by M and M_o . Then because $\text{RSS}_{M'} \leq \text{RSS}_M$,

$$P_y\left(\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right) \leq P_y\left(\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_{M'}}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right) \leq P_y(\chi_{j+l}^2 > \chi_{n-j-l}^2(a_n - 1)),$$

and for any nonnegative $s \in \mathbb{R}$,

$$(19) \quad \begin{aligned} P_y\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > a_n\right) &\leq P_y(\{\chi_{j+l}^2 > s(a_n - 1)\} \cap \{\chi_{n-j-l}^2 \geq s\}) \\ &\quad + P_y(\{\chi_{j+l}^2 > \chi_{n-j-l}^2(a_n - 1)\} \cap \{\chi_{n-j-l}^2 < s\}) \\ &\leq P_y(\chi_{j+l}^2 > s(a_n - 1)) + P_y(\chi_{n-j-l}^2 < s). \end{aligned}$$

Consider each of these last two terms in turn. For the first term, apply the Chernoff bound, and evaluate the moment generating function for the χ_{j+l}^2 distribution at $1/4$. That gives

$$P_y(\chi_{j+l}^2 > s(a_n - 1)) \leq 2^{\frac{j+l}{2}} e^{-\frac{s(a_n-1)}{4}} \leq e^{-\frac{s(a_n-1)}{4} + \frac{j+l}{2}}.$$

For the second term in (19), write out the expression to evaluate the probability explicitly, and then apply the simple bound $e^{-x} \leq 1$ for all $x \geq 0$. Noting that $s > 0$,

$$P_y(\chi_{n-j-l}^2 < s) \leq \frac{1}{2^{\frac{n-j-l}{2}} \Gamma(\frac{n-j-l}{2})} \frac{s^{\frac{n-j-l}{2}}}{\frac{n-j-l}{2}} \leq \frac{(e \cdot s)^{\frac{n-j-l}{2}} \cdot 2^{-\frac{n-j-l}{2}}}{\sqrt{2\pi} (\frac{n-j-l}{2})^{\frac{n-j-l}{2} + \frac{1}{2}}},$$

where the last inequality follows from the well-known Sterling lower bound on the gamma function, $\Gamma(x) \geq \sqrt{2\pi} x^{x-\frac{1}{2}} e^{-x}$ for $x > 0$. It is clear from the last expression that for the probability to vanish, $e \cdot s$ must grow no faster than $\frac{n-j-l}{2}$. Accordingly, choosing $s = \frac{n-j-l}{e^2}$ gives

$$P_y\left(\chi_{n-j-l}^2 < \frac{n-j-l}{e^2}\right) \leq \frac{e^{-\frac{n-j-l}{2}}}{\sqrt{\pi(n-j-l)}}.$$

Combining the two bounds for (19) now yields

$$P_y\left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M} > a_n\right) \leq e^{-\left(\frac{n-j-l}{e^2}\right)\frac{a_n-1}{4} + \frac{j+l}{2}} + \frac{e^{-\frac{n-j-l}{2}}}{\sqrt{\pi(n-j-l)}}.$$

It only remains to choose the smallest a_n such that the first term in the bound vanishes exponentially fast so that the cumulative probability will vanish in probability over all subsets M with $|M| \leq n^\alpha$. Accordingly, it should become apparent

shortly that a good choice is

$$(20) \quad a_n = 1 + \frac{8e^2(n^\alpha + j \log(p))}{n - j - 1}.$$

The probability that $(\frac{RSS_{M_o}}{RSS_M})^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}$ is satisfied for some M with $|M_o| < |M| \leq n^\alpha$ is

$$\begin{aligned} & P_y \left(\bigcup_{j=|M_o|+1}^{n^\alpha} \bigcup_{M:|M|=j} \left\{ \left(\frac{RSS_{M_o}}{RSS_M} \right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}} \right\} \right) \\ & \leq \sum_{j=|M_o|+1}^{n^\alpha} \binom{p}{j} \max_{M:|M|=j} P_y \left(\left(\frac{RSS_{M_o}}{RSS_M} \right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}} \right). \end{aligned}$$

Thus, bounding the binomial coefficient as in (18), and substituting (20) for a_n yields

$$\begin{aligned} & P_y \left(\bigcup_{j=|M_o|+1}^{n^\alpha} \bigcup_{M:|M|=j} \left\{ \left(\frac{RSS_{M_o}}{RSS_M} \right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}} \right\} \right) \\ & \leq \sum_{j=|M_o|+1}^{n^\alpha} e^{-b_n[2n^\alpha - \frac{j}{2b_n} + (2 - \frac{1}{b_n} - \frac{1}{2b_n \log(p)})j \log(p)]} + \frac{e^{-\frac{n-j-|M_o|}{2} + j \log(p)}}{\sqrt{\pi(n - j - |M_o|)}} \\ & \leq n^\alpha e^{-b_n[2n^\alpha - \frac{|M_o|}{2b_n} + (2 - \frac{1}{b_n} - \frac{1}{2b_n \log(p)})(|M_o|+1) \log(p)]} \\ & \quad + \frac{n^\alpha e^{-\frac{n-n^\alpha-|M_o|}{2} + n^\alpha \log(p)}}{\sqrt{\pi(n - n^\alpha - |M_o|)}}, \end{aligned}$$

where $b_n := \frac{n-n^\alpha-|M_o|}{n-|M_o|-2} \rightarrow 1$. Note that this bound vanishes by Condition 3.2. Therefore, with probability exceeding one minus the above bound,

$$\left(\frac{RSS_{M_o}}{RSS_M} \right)^{\frac{n-j-1}{2}} \leq \left(1 + \frac{4e^2(n^\alpha + j \log(p))}{\frac{n-j-1}{2}} \right)^{\frac{n-j-1}{2}} \leq e^{4e^2(n^\alpha + j \log(p))},$$

uniformly over all M such that $|M_o| < |M| \leq n^\alpha$. \square

PROOF OF THEOREM 3.9. Without loss of generality suppose $\sigma_{M_o}^0 = 1$, where $\sigma_{M_o}^0$ is the true but unknown error standard deviation. For $j \in \{1, \dots, p\}$, define the following classes of subsets $\mathcal{M}_j := \{M \neq M_o : |M| = j\}$. Recall that $\hat{\sigma}_M^2 := RSS_M / (n - |M|)$. It will first be shown that for any subset of covariates $M \neq M_o$ the ratio $\frac{r(M|Y)}{r(M_o|Y)}$ vanishes in probability for large n and p . Accordingly, for

any $M \in \mathcal{M}_j$,

$$\begin{aligned} \frac{r(M|Y)}{r(M_o|Y)} &= \pi^{\frac{j-|M_o|}{2}} \frac{\Gamma(\frac{n-j}{2})}{\Gamma(\frac{n-|M_o|}{2})} \frac{\text{RSS}_{M_o}^{\frac{n-|M_o|-1}{2}}}{\text{RSS}_M^{\frac{n-j-1}{2}}} \frac{\mathbb{E}(h(\beta_M))}{\mathbb{E}(h(\beta_{M_o}))} \\ &= \pi^{\frac{j-|M_o|}{2}} \frac{\Gamma(\frac{n-j}{2})}{\Gamma(\frac{n-|M_o|}{2})} \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M}\right)^{\frac{n-j-1}{2}} \text{RSS}_{M_o}^{\frac{j-|M_o|}{2}} \frac{\mathbb{E}(h(\beta_M))}{\mathbb{E}(h(\beta_{M_o}))}. \end{aligned}$$

Before proceeding with the rest of the proof, the following notation is needed. V_1 and V_2 are as stated in Lemma A.1. As in Theorem 3.8, define

$$V_3 := \frac{3p_{M_o} \sigma_M \sqrt{\Lambda_{M_o}}}{\sqrt{\varepsilon \pi}} e^{-\frac{\varepsilon}{9\sigma_M^2 \Lambda_{M_o}}},$$

and corresponding to Theorem 3.7, define

$$V_4 := \sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{3|M| \sigma_M \sqrt{\Lambda_M}}{\sqrt{\varepsilon \pi}} e^{-\frac{\varepsilon}{9\sigma_M^2 \Lambda_M}} \leq \sum_{j=1}^{n^\alpha} \max_{\mathcal{M}_j} \frac{3|M| \sigma_M \sqrt{\Lambda_M}}{\sqrt{\varepsilon \pi} e^{\frac{\varepsilon}{9\sigma_M^2 \Lambda_M} + j \log(p)}}$$

by bounding the binomial coefficient as in (18). Note that V_4 then vanishes by Condition 3.4. Further, recall that $\text{RSS}_{M_o} \sim \chi_{n-|M_o|}^2$, so by the Chernoff bound (evaluating the moment generating function at $1/4$),

$$P_Y(\chi_{n-|M_o|}^2 > 3(n - |M_o|)) \leq \underbrace{e^{-\frac{n-|M_o|}{4}}}_{=: V_5}.$$

With this notation, it is now possible to account for all of the uncertainty due to Y .

Accordingly, by Theorem 3.8, with probability exceeding $1 - V_3 - V_5$,

$$(21) \quad \frac{r(M|Y)}{r(M_o|Y)} \leq \frac{\Gamma(\frac{n-j}{2})}{\Gamma(\frac{n-|M_o|}{2})} \left(\frac{\text{RSS}_{M_o}}{\text{RSS}_M}\right)^{\frac{n-j-1}{2}} \frac{(3\pi(n - |M_o|))^{\frac{j-|M_o|}{2}} \mathbb{E}(h(\beta_M))}{1 - g_1(M_o, n, p)},$$

where

$$g_1(M, n, p) := \frac{2^{\frac{3}{2}} 3 j \hat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n-j})} e^{-\frac{\varepsilon}{18 \Lambda_M \hat{\sigma}_M^2}}.$$

Further, fix $A_1 \in (0, 1)$, and by Condition 3.4 choose n and p sufficiently large so that $g_1(M_o, n, p) < A_1$.

Suppose $j \leq |M_o|$. As in Jameson [12], the ratio of gamma functions can be bounded by

$$\frac{\Gamma(\frac{n-j}{2})}{\Gamma(\frac{n-|M_o|}{2})} = \frac{\Gamma(\frac{n-|M_o|}{2} + \frac{|M_o|-j}{2})}{\Gamma(\frac{n-|M_o|}{2})} \leq \binom{n - |M_o|}{2} \binom{n - j}{2}^{\frac{|M_o|-j}{2} - 1}.$$

Applying Lemma A.1 to bound the ratio of residual sums of squares, and bounding the expectation by 1, with probability exceeding $1 - V_1 - V_3 - V_5$, (21) implies

$$(22) \quad \frac{r(M|Y)}{r(M_o|Y)} \leq \frac{\binom{\frac{n-|M_o|}{2}}{\frac{j-|M_o|}{2}} (3\pi(n - |M_o|))^{\frac{j-|M_o|}{2}}}{\binom{\frac{n-j}{2}-\frac{|M_o|-j}{2}+1} (1 - A_1)e^{2\log(n)^\gamma \log(p)}} \leq \frac{e^{-2\log(n)^\gamma \log(p)}}{1 - A_1},$$

where the last inequality follows for all $n \geq \frac{|M_o|}{1-\frac{1}{6\pi}}$. Since $n \gg |M_o|$, assume without loss of generality that n is sufficiently large.

Conversely, suppose $|M_o| < j \leq n^\alpha$. In this setting, as in Jameson [12], the ratio of gamma functions can be bounded by

$$\left(\frac{\Gamma(\frac{n-j}{2})}{\Gamma(\frac{n-|M_o|}{2})}\right)^{-1} = \frac{\Gamma(\frac{n-|M_o|}{2})}{\Gamma(\frac{n-j}{2})} = \frac{\Gamma(\frac{n-j}{2} + \frac{j-|M_o|}{2})}{\Gamma(\frac{n-j}{2})} \geq \left(\frac{n-j}{2} - 1\right)^{\frac{j-|M_o|}{2}}.$$

Applying Theorem 3.7 to bound the expectation, and applying Lemma A.1 to bound the ratio of residual sums of squares, with probability exceeding $1 - V_2 - V_3 - V_4 - V_5$, (21) implies

$$(23) \quad \begin{aligned} \frac{r(M|Y)}{r(M_o|Y)} &\leq \frac{e^{4e^2(n^\alpha + j \log(p))}}{\binom{\frac{n-j}{2} - 1}^{\frac{j-|M_o|}{2}}} (3\pi(n - |M_o|))^{\frac{j-|M_o|}{2}} \frac{g_1(M, n, p)}{1 - A_1} \\ &\leq \frac{2^{\frac{3}{2}} 3j \hat{\sigma}_M \sqrt{\Lambda_M} e^{-\frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2} + 4e^2(n^\alpha + j \log(p)) + \frac{j-|M_o|}{2} \log(\frac{6\pi}{1-\frac{6\pi}{n}+2})}}{\sqrt{\pi \varepsilon} (1 - \frac{1}{n-j})(1 - A_1)}. \end{aligned}$$

Notice that Theorem 3.7 is being applied here with n^α in place of $c_1 n$. This can be done without loss of generality because n^α grows slower than $c_1 n$, for any choices of $\alpha, c_1 \in (0, 1)$.

It can now be shown that $\sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{r(M|Y)}{r(M_o|Y)}$ vanishes in probability for large n and p . Apply the bounds in (22) and (23) in the following argument:

$$\sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{r(M|Y)}{r(M_o|Y)} \leq \underbrace{\sum_{j=1}^{|M_o|} \binom{p}{j} \max_{M \in \mathcal{M}_j} \frac{r(M|Y)}{r(M_o|Y)}}_{=:S_1} + \underbrace{\sum_{j=|M_o|+1}^{n^\alpha} \binom{p}{j} \max_{|M|=j} \frac{r(M|Y)}{r(M_o|Y)}}_{=:S_2}.$$

Consider S_1 and S_2 separately.

With probability exceeding $1 - V_1 - V_3 - V_5$,

$$S_1 \leq \frac{1}{1 - A_1} \sum_{j=1}^{|M_o|} e^{-2\log(n)^\gamma \log(p) + j \log(p)} \leq \frac{|M_o| e^{-(2\log(n)^\gamma - |M_o|) \log(p)}}{1 - A_1},$$

by bounding the binomial coefficient as in (18), and with probability exceeding $1 - V_2 - V_3 - V_4 - V_5$,

$$\begin{aligned}
 S_2 &\leq A_2 \sum_{j=|M_o|+1}^{n^\alpha} \frac{\frac{j}{1-\frac{1}{n-j}} \max_{|M|=j} \left(\frac{\hat{\sigma}_M^2 \Lambda_M}{\varepsilon}\right)^{\frac{1}{2}}}{e^{\min_{|M|=j} \left\{ \frac{\varepsilon}{18\Lambda_M \hat{\sigma}_M^2} - 4e^2(n^\alpha + j \log(p)) - \frac{j-|M_o|}{2} \log\left(\frac{6\pi}{1-\frac{n^\alpha+2}{n}}\right) - j \log(p) \right\}}} \\
 &\leq \frac{A_2}{n^{1-\alpha}} \cdot \max_{|M_o| < |M| \leq n^\alpha} \left(\frac{\hat{\sigma}_M^2 \Lambda_M}{\varepsilon}\right)^{\frac{1}{2}}
 \end{aligned}$$

for some positive constant A_2 . The last inequality follows by Condition 3.4.

Thus, for sufficiently large n and p , with probability exceeding $1 - V_1 - V_2 - V_3 - V_4 - V_5$,

$$\sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{r(M|Y)}{r(M_o|Y)} \leq \frac{|M_o| e^{-(2 \log(n)^\gamma - |M_o|) \log(p)}}{1 - A_1} + \frac{\max_{|M_o| < |M| \leq n^\alpha} \left(\frac{\hat{\sigma}_M^2 \Lambda_M}{\varepsilon}\right)^{\frac{1}{2}}}{A_2^{-1} n^{1-\alpha}},$$

which by Conditions 3.1–3.4 vanishes for large n and p . The proof is now complete by noticing that

$$r(M_o|Y) = \frac{r(M_o|Y)}{\sum_{j=1}^{n^\alpha} \sum_{M:|M|=j} r(M|Y)} = \frac{1}{1 + \sum_{j=1}^{n^\alpha} \sum_{\mathcal{M}_j} \frac{r(M|Y)}{r(M_o|Y)}}. \quad \square$$

REFERENCES

- [1] ANDRIEU, C. and ROBERTS, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.* **37** 697–725. [MR2502648](#)
- [2] BEAUMONT, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics* **164** 1139–1160.
- [3] BERGER, J. O. and PERICCHI, L. R. (2001). Objective Bayesian methods for model selection: Introduction and comparison. In *Model Selection. Institute of Mathematical Statistics Lecture Notes—Monograph Series* **38** 135–207. IMS, Beachwood, OH. [MR2000753](#)
- [4] BERK, R. A. (2008). *Statistical Learning from a Regression Perspective*. Springer, New York. [MR2722293](#)
- [5] BERTSIMAS, D., KING, A. and MAZUMDER, R. (2016). Best subset selection via a modern optimization lens. *Ann. Statist.* **44** 813–852. [MR3476618](#)
- [6] BONDELL, H. D. and REICH, B. J. (2012). Consistent high-dimensional Bayesian variable selection via penalized credible regions. *J. Amer. Statist. Assoc.* **107** 1610–1624. [MR3036420](#)
- [7] BREHENY, P. and HUANG, J. (2011). Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.* **5** 232–253. [MR2810396](#)
- [8] CANDÈS, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35** 2313–2351. [MR2382644](#)
- [9] FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- [10] GHOSH, J. and GHATTAS, A. E. (2015). Bayesian variable selection under collinearity. *Amer. Statist.* **69** 165–173. [MR3391636](#)

- [11] HANNIG, J., IYER, H., LAI, R. C. S. and LEE, T. C. M. (2016). Generalized fiducial inference: A review and new results. *J. Amer. Statist. Assoc.* **111** 1346–1361. [MR3561954](#)
- [12] JAMESON, G. J. O. (2013). Inequalities for gamma function ratios. *Amer. Math. Monthly* **120** 936–940. [MR3139572](#)
- [13] JOHNSON, V. E. and ROSSELL, D. (2012). Bayesian model selection in high-dimensional settings. *J. Amer. Statist. Assoc.* **107** 649–660. [MR2980074](#)
- [14] LAI, R. C. S., HANNIG, J. and LEE, T. C. M. (2015). Generalized fiducial inference for ultrahigh-dimensional regression. *J. Amer. Statist. Assoc.* **110** 760–772. [MR3367262](#)
- [15] LUO, S. and CHEN, Z. (2013). Extended BIC for linear regression models with diverging number of relevant features and high or ultra-high feature spaces. *J. Statist. Plann. Inference* **143** 494–504. [MR2995110](#)
- [16] NARISSETTY, N. N. and HE, X. (2014). Bayesian variable selection with shrinking and diffusing priors. *Ann. Statist.* **42** 789–817. [MR3210987](#)
- [17] PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R. et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12** 2825–2830. [MR2854348](#)
- [18] ROČKOVÁ, V. and GEORGE, E. I. (2018). The Spike-and-Slab LASSO. *J. Amer. Statist. Assoc.* **113** 431–444. [MR3803476](#)
- [19] ROSSELL, D. and TELESKA, D. (2017). Nonlocal priors for high-dimensional estimation. *J. Amer. Statist. Assoc.* **112** 254–265. [MR3646569](#)
- [20] SHIN, M., BHATTACHARYA, A. and JOHNSON, V. E. (2018). Scalable Bayesian variable selection using nonlocal prior densities in ultrahigh-dimensional settings. *Statist. Sinica* **28** 1053–1078. [MR3791100](#)
- [21] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- [22] ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the LASSO selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. [MR2435448](#)

DEPARTMENT OF STATISTICS AND
OPERATIONS RESEARCH
UNIVERSITY OF NORTH CAROLINA AT CHAPEL HILL
CHAPEL HILL, NORTH CAROLINA 27599
USA
E-MAIL: jpwill@live.unc.edu
jan.hannig@unc.edu