

The EAS approach to variable selection for multivariate response data in high-dimensional settings

Salil Koner¹ and Jonathan P. Williams²

¹*Duke University, Durham, NC*
e-mail: Salil.Koner@duke.edu

²*North Carolina State University, Raleigh, NC*
e-mail: jwilli27@ncsu.edu

Abstract: In this paper, we develop an *epsilon admissible subsets* (EAS) model selection approach for performing group variable selection in the high-dimensional multivariate regression setting. This EAS strategy is designed to estimate a posterior-like, generalized fiducial distribution over a parsimonious class of models in the setting of correlated predictors and/or in the absence of a sparsity assumption. The effectiveness of our approach, to this end, is demonstrated empirically in simulation studies, and is compared to other state-of-the-art model/variable selection procedures. Furthermore, assuming a matrix-Normal linear model we show that the EAS strategy achieves *strong model selection consistency* in the high-dimensional setting if there does exist a sparse, true data generating set of predictors. In contrast to Bayesian approaches for model selection, our generalized fiducial approach completely avoids the problem of simultaneously having to specify arbitrary prior distributions for model parameters and penalize model complexity; our approach allows for inference directly on the model complexity. Implementation of the method is illustrated through yeast data to identify significant cell-cycle regulating transcription factors.

MSC2020 subject classifications: Primary 62H12.

Keywords and phrases: Generalized fiducial inference, model selection, regularized regression, asymptotic consistency.

Received June 2022.

1. Introduction

With the advent of modern data collection technologies in many real-life applications, multiple responses are simultaneously collected that are characterized by a set of explanatory variables. Data of this structure falls under the scope of multivariate regression. Examples arise in chemometrics [18], genome-wide association studies (GWAS) [7], etc. More often than not, the number of predictors p is much larger than the number of observed multivariate response vectors n . Parsimoniously modeling the variability in the response without overfitting is necessary to enhance prediction accuracy. For example, in GWAS, identification of key genetic markers out of millions that are associated with a univariate or multivariate phenotype is of scientific interest [50]. Model/variable selection is

a statistical framework that has been widely popular in this context. A naive way to approach this multivariate problem is to model the components of the response as separate univariate regressions on the predictors, and to employ existing selection techniques available in the univariate setting. However, since the multiple responses for each subject are often correlated, prediction error can be minimized substantially if one uses the inherent association between the responses effectively [8]. Finding the best model in the context of multivariate linear regression (MLR) without ignoring the correlation between the responses, especially in a high-dimensional setting is a challenging task and has received much attention in the literature over the last decade.

We develop an EAS approach for *group* variable selection in high-dimensional MLR settings. The EAS procedure was originally developed for high-dimensional univariate regression settings of [53], and has been extended to the vector autoregression setting in [54]. However, an EAS procedure has not been constructed for the multivariate regression setting, nor has there been built a group selection mechanism for EAS selection. In contrast to Bayesian model selection approaches, we consider a generalized fiducial (GF) inference approach [24] that explicitly estimates the GF distribution over the class of all subsets of predictors; whereas frequentist and most Bayesian approaches exclusively focus on coefficient estimation to perform variable selection. In a high-dimensional setting, it is very problematic for a variable selection procedure to over-rely on the magnitude of the estimated regression coefficients because they lack identifiability and are numerically unstable. Moreover, most variable selection procedures for multivariate regression do *not* account for the correlation structure of the multivariate responses. Our multivariate group EAS procedure is designed to inherently accommodate the arbitrary covariance structure of the response. As explained in [28], accounting for the correlation is important because, for example, in the case of the least absolute shrinkage and selection operator (LASSO) estimator the shrinkage criterion is affected by the magnitude and sign of the correlation between the responses.

Mathematically, under a standard sparsity and Gaussian errors assumption with a general covariance structure we prove that our proposed EAS procedure achieves *strong model selection* consistency, as defined in [36]. That is, we show that over the class of all ϵ -admissible subsets of the predictors/groups, the GF probability of a true sparse model converges to one in probability as the sample size goes to infinity, and the number of predictors/groups is allowed to grow sub-exponentially as a function of the sample size. Additionally, as a next paper in the series of papers to investigate EAS model selection strategies in various settings, the algorithm we propose improves on the computational efficiency and stability of the algorithms proposed in [53, 54]. We provide user-friendly R software to implement our EAS procedure, available at <https://github.com/SalilKoner/EAS>.

Next, in various simulation scenarios, we demonstrate that our EAS procedure is either competitive with or outperforms the state-of-the-art Bayesian or frequentist approaches, based on various metrics such as prediction error, misclassification rate, and average proportion of correct model selections. Moreover,

consistent with our theory, our EAS method does an excellent job in assigning a very high probability mass to the true model compared to the other Bayesian methods.

There are four distinct elements of novelty of our EAS procedure for the MLR setting compared to the univariate case [53]. To start with the methodological novelties: 1) Our EAS method is designed to inherently accommodate an arbitrary covariance structure of the response in performing variable selection. 2) The second added difficulty, specific to GF inference, is that the so-called GF “Jacobian term” for the MLR model (derived in Appendix A) is non-standard and exerts substantial influence on the resulting GF model probabilities that we construct, whereas the Jacobian term in the univariate linear regression setting is concise and involves components readily relatable to the likelihood function. 3) The theoretical novelties, primarily imparted by accounting for an arbitrary covariance structure, constitute derivation of non-asymptotic concentration bounds on ratios of determinants (or eigenvalues) of error covariance matrices for different models (i.e., subsets of covariates) in the multivariate setting, for which much less standard theory exists. 4) The key computational improvement of our EAS method is that the user does not need to scale the tuning parameter ϵ as a function of the sample size to select the best model, in contrast to previous EAS developments such as [53].

Prior choice/specification in contemporary Bayesian approaches are typically not chosen because they reflect true prior knowledge/beliefs, but they are tailored to simplify computational complexities and/or achieve desirable large sample/frequentist properties. While this practice is pragmatic, it is a violation of fundamental Bayesian principles. In contrast, GF inference is an equally principled framework that has an appeal to objective Bayesian perspectives, but it does not suffer from the arbitrary choice of prior specification. The GF approach is to solve an inverse problem resulting in parameter values most consistent in reconciling the observed data with random draws from the distribution of the auxiliary variable in the data generating equation. Holding the data fixed, these parameter values then inherit a probability distribution via the distribution of the auxiliary variable. It can be argued that this procedure effectively has implicit prior knowledge built in, but the subjectivist Bayesian approach (assuming a likelihood function) goes a step further by imposing/requiring additional prior knowledge (in the form of a prior distribution specification) that is exogenous to the data generating model. Moreover, it has been shown that GF inference exhibits large sample Bernstein-von Mises type properties that guarantee the nominal coverage of credible sets, similar to such theory for Bayesian posteriors [41]. See [24] for a full introduction of GF inference.

Standard model selection techniques for univariate linear regression via Mallows’s C_p and other types of information criterion have been extended to multivariate regression; see [42, 20] and the references therein. Since the inception of LASSO [43], penalized methods introducing sparsity in the regression coefficients have engulfed the MLR literature. Notable contributions are the simultaneous variable selection using L_∞ SVS from [45] and L_2 SVS from [40]; the remMap procedure in [37] employs an elastic net type penalty to identify

master predictors. Estimation of the regression coefficients taking into account the correlation between the responses was pioneered in [39], which was later extended in [28]. Assuming that the groups are known, a multivariate sparse group LASSO strategy is proposed in [29] to impose the group structure, which was later augmented in [55] for simultaneous covariance estimation. Very recently, de-sparsified/de-biased LASSO was developed in [14, 4] to overcome the bias induced by the penalty, by extending the existing techniques developed for the univariate setting [46]. From a Bayesian perspective, a stochastic search variable selection (SSVS) procedure for MLR is employed in [9, 10]. The multivariate Bayesian group LASSO using a spike-and-slab prior is developed in [31] to perform the variable selection. To mitigate the computational issues of spike-and-slab priors for the large p scenario, continuous global-local shrinkage priors are introduced in [1]. Recently, an expectation-maximization based maximum a posteriori (MAP) estimation procedure was formulated for fast simultaneous variable and covariance selection using continuous shrinkage priors in [17].

So far the NP-hard problem of best subset selection has been seemingly conveniently handled by assuming sparsity in the true data-generating model. However, in a high-dimensional setting, especially when there is a strong degree of collinearity amongst the predictors, there may not be a unique model that fits the data best, and so the concept of a *true model* is somewhat vague. Moreover, the typical ℓ_1 and ℓ_2 regularization methods shrink coefficients to zero only based on their magnitude, which is again unreliable in the presence of multicollinearity. Our EAS procedure provides a fresh perspective on variable selection in the MLR setting by defining an admissibility condition for candidate models. The admissibility criterion relates to the idea that any candidate model, as defined by a set of predictors, is redundant if there exists a subset model that explains the variation in the response as well as the candidate model. Thus, while we prove important mathematical properties of the EAS approach under a sparsity assumption, in finite sample data analyses the key functionality of our EAS approach is not to determine the model that is necessarily the true set of predictors but to identify a parsimonious model that explains the data as well as the true model, if it exists. This criterion is meaningful even in the absence of a sparsity assumption. We refer to the criterion as ϵ -admissibility, defined in Section 2. The key characteristic of the EAS-based GF distribution is that it assigns very negligible probability to models that fail the ϵ -admissibility criterion, and in doing so significantly reduces the class of candidate models to choose from. This is the intuition for why, assuming sparsity, the procedure achieves strong model selection consistency.

An advantage of our GF-based variable selection procedure over the frequentist counterparts is that it provides an estimate of the model probabilities, derived from the posterior-like GF distribution of the model parameters. Even many recent standard Bayesian approaches, such as Multivariate Bayesian model with Shrinkage Priors (MBSP) from [1], multivariate spike-and-slab prior (mSSL) from [17], and Spike-and-Slab Group LASSO (SSGL) from [3] are not developed to compute relative model probabilities; rather they are designed to estimate the MAP probability model parameters. A notable exception is the

Bayesian spike-and-slab prior group LASSO from [31] that is capable of providing relative model probabilities, but is unfortunately not suitable for high-dimensional settings. In a finite sample scenario, relative model probabilities are useful because they give a certain degree of confidence in choosing one model over another, and they reflect a useful discrimination between competing models, especially in the high-dimensional setting. Moreover, the ϵ -admissibility criterion in our EAS framework takes into account the covariance structure of the multivariate response. In a multi-response setting, the noise associated with a particular component of the response may be significantly higher than the other components. Many of the variable selection procedures such as remMap [37], L_2 SVS [40], or reduced rank regression [13, 48], neither consider this difference in noise levels nor the intra-dependence in the multivariate response.

The rest of the paper is organized as follows. In Section 2 we layout the EAS methodology and highlight differences from other EAS approaches. Next, the computational algorithm to implement the method is presented in Section 3. A few essential non-asymptotic results characterizing the meaningfulness of the EAS procedure along with the main consistency result are stated in Section 4. The proof of all the results are relegated to the Section B of the Appendix. Finite sample numerical results covering both $n > p$ and $p > n$ cases are illustrated in Section 5. Section 6 provides an illustration of the practical application of the procedure using yeast cell cycle data. Computer codes to reproduce all empirical results are available at <https://github.com/SalilKoner/EAS>.

Notations

Throughout the course of the paper we will use the following notations. For an vector $\mathbf{a} \in \mathbb{R}^n$, $\|\mathbf{a}\| =: \sqrt{\sum_{i=1}^n a_i^2}$ denotes vector 2-norm; for any matrix \mathbf{A} , $\|\mathbf{A}\|$ refers to the spectral norm (i.e., $\|\mathbf{A}\| := \sup_{\mathbf{x}, \|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|$); $\|\mathbf{A}\|_F$ denotes the Frobenius norm; and $\|\mathbf{A}\|_{\max} := \max_{i,j} |a_{ij}|$ denotes the max norm. For any symmetric matrix \mathbf{A} , $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denotes the minimum and maximum eigenvalues, respectively, of the matrix \mathbf{A} . For an $m \times r$ random matrix \mathbf{X} , $\mathbf{X} \sim \text{Matrix-Normal}_{m,r}(\mathbf{M}, \mathbf{U}, \mathbf{V})$ is equivalent to $\text{vec}(\mathbf{M}) \sim N_{m,r}(\text{vec}(\mathbf{M}), \mathbf{V} \otimes \mathbf{U})$ [see 22, definition 2.2.1]. For a $m \times r$ random matrix \mathbf{X} , $\mathbf{X} \sim T_{m,r}(\nu, \mathbf{M}, \mathbf{U}, \mathbf{V})$ means \mathbf{X} follows a matrix t -distribution with mean \mathbf{M} , scale matrices \mathbf{U} , \mathbf{V} , and degrees of freedom ν [see 22, definition 4.2.1]. For any set M , $|M|$ denotes the cardinality of M . For any event A , $\mathbf{I}(A)$ denotes the indicator function that the event happens. Lastly, for two sequences a_n and b_n , $a_n = o(b_n)$ means $\lim_{n \rightarrow \infty} a_n/b_n = 0$.

2. Methodology

In MLR, n pairs of examples $(\mathbf{X}_i, \mathbf{Y}_i)$ are observed, where $\mathbf{Y}_i := (Y_{i1}, \dots, Y_{iq})^\top$ is the q -dimensional multivariate response vector for i -th subject, and $\mathbf{X}_i := (X_{1i}, \dots, X_{pi})^\top$ contains the values of p predictor variables that are presumed

to be associated with the response. The response is expressed as the linear model,

$$\mathbf{Y}_i = \sum_{j=1}^p \mathbf{B}_j X_{ji} + \mathbf{A}\mathbf{U}_i,$$

where $\mathbf{B}_j := (b_{1j}, \dots, b_{qj})^\top$ is a q -dimensional regression coefficient vector where $b_{\ell j}$ captures the effect of j -th predictor on ℓ -th coordinate of the multivariate response, \mathbf{A} is $q \times q$ matrix, and $\mathbf{U}_i \in \mathbb{R}^{q \times 1}$ is the i -th error vector with mean $\mathbf{0}$ and $\text{Var}(\mathbf{U}_i) = \mathbf{I}_q$, so that $\text{Var}(\mathbf{Y}_i) = \mathbf{A}\mathbf{A}^\top$. Further denoting $\mathbf{Y} := [\mathbf{Y}_1, \dots, \mathbf{Y}_n] \in \mathbb{R}^{q \times n}$ as the horizontal column stacked response, and $\mathbf{U} := [\mathbf{U}_1, \dots, \mathbf{U}_n]$ as the corresponding $q \times n$ dimensional error matrix, the multivariate regression model with a sample of size n is summarized as,

$$\mathbf{Y} = \left[\sum_{j=1}^p \mathbf{B}_j X_{j1}, \dots, \sum_{j=1}^p \mathbf{B}_j X_{jn} \right] + \mathbf{A}\mathbf{U} = \mathbf{B}\mathbf{X} + \mathbf{A}\mathbf{U}, \quad (1)$$

where $\mathbf{B} := [\mathbf{B}_1, \dots, \mathbf{B}_p] \in \mathbb{R}^{q \times p}$ is the coefficient matrix and $\mathbf{X} := [\mathbf{X}_1, \dots, \mathbf{X}_n] \in \mathbb{R}^{p \times n}$ is the design matrix.

In the context of variable selection, for any index set $M \subseteq \{1, \dots, p\}$, let \mathbf{X}_M denote the matrix with *rows* comprised of the rows of \mathbf{X} corresponding to the indices in M . In subsetting the rows of \mathbf{X} for variable selection, the *columns* of \mathbf{B} must be subset to only those corresponding to the indices in M ; take $\mathbf{B}_{\cdot M}$ to be the column-subsetted coefficient matrix. The subscript (M) , as in the $q \times q$ matrix $\mathbf{A}_{(M)}$, simply denotes an association with the index set M , rather than any subsetting of the rows/columns. Accordingly, conditional on index set/model M , expression (1) reduces to,

$$\mathbf{Y} = \mathbf{B}_{\cdot M} \mathbf{X}_M + \mathbf{A}_{(M)} \mathbf{U}. \quad (2)$$

Observe that variable selection in this multivariate model is in fact a group selection problem because the active columns (i.e., the groups) of the coefficient matrix \mathbf{B} are being selected. This fact establishes the need for the a group selection mechanism within the variable selection procedure. Nonetheless, the EAS variable selection procedure that we develop seamlessly accommodates the additional problem of selecting among known/posited groups of predictors, rather than the power set of the predictors $1, \dots, p$ (i.e., the *natural grouping*). Our presentation will focus on EAS methodology for solving the problem of grouped variable selection for MLR under the natural grouping, but the methodology and theoretical results cover (as a sub-case) the simpler case when a class of predictor groups are known. In that case, simply restrict the class of candidate models that the algorithm is allowed to choose from.

In the case of Gaussian error, we introduce the notion of a *true* data generating model as the assumption that,

$$\mathbf{Y} \sim \text{Matrix-Normal}_{q,n} \left(\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o}, \mathbf{V}_{(M_o)}^0, \mathbf{I}_n \right), \quad (3)$$

for some fixed (but unknown) $M_o \subseteq \{1, 2, \dots, p\}$, and some fixed (but unknown) parameter matrices $\mathbf{B}_{\cdot M_o}^0$ and $\mathbf{V}_{(M_o)}^0 = \mathbf{A}_{(M_o)}^0 \mathbf{A}_{(M_o)}^{0\top}$, for some positive definite matrix $\mathbf{A}_{(M_o)}^0$. Note that in the case of the covariance matrix $\mathbf{V}_{(M_o)}^0$, the subscript M_o simply denotes the association with the index set M_o but it is *not* constructed by subsetting some more general matrix \mathbf{V}^0 . Analogous to [53], the index ‘o’ in M_o is in reference to the term ‘oracle’, and the superscript ‘0’ emphasizes that the quantities $\mathbf{B}_{\cdot M_o}^0$ and $\mathbf{V}_{(M_o)}^0$ are fixed quantities (in contrast to their GF/Bayesian-like random variable counterparts $\mathbf{B}_{\cdot M_o}$ and $\mathbf{V}_{(M_o)}$, to be introduced shortly). The matrix normal notation in (3) is a compact way of saying that the multivariate responses $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are independent and identically distributed multivariate normal random vectors with mean $\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o}$ and $q \times q$ covariance matrix $\mathbf{V}_{(M_o)}^0$.

The objective of our paper is to develop a methodology that identifies a *non-redundant* (i.e., ϵ -admissible) set M_o (the true model or possibly a related sub-model of the true model) out of the 2^p candidate sets in the power set of $\{1, 2, \dots, p\}$. In the special case that the true model is sparse, the objective is to show that the method will identify *the* true model M_o , as the sample size is taken to infinity. The notion of ϵ -admissibility is defined in Definition 2.1. To build up to this definition, the following model-based perspective is required.

For any index set $M \subseteq \{1, \dots, p\}$, we assume the conditional distribution,

$$\mathbf{Y} | \mathbf{B}_{\cdot M}, \mathbf{V}_{(M)} \sim \text{Matrix-Normal}_{q,n}(\mathbf{B}_{\cdot M} \mathbf{X}_{M\cdot}, \mathbf{V}_{(M)}, \mathbf{I}_n), \quad (4)$$

where $\mathbf{B}_{\cdot M}$ and $\mathbf{V}_{(M)} = \mathbf{A}_{(M)} \mathbf{A}_{(M)}^\top$ are random matrices that reflect the uncertainty in not knowing the true data generating model nor true values of its parameter matrices. Under the Gaussian error assumption, the quantities $\mathbf{B}_{\cdot M}$ and $\mathbf{V}_{(M)}$ are expected to be centered, respectively, around the least squares estimator $\widehat{\mathbf{B}}_{\cdot M} := \mathbf{Y} \mathbf{X}_M^\top (\mathbf{X}_M \mathbf{X}_M^\top)^{-1}$ and the restricted maximum likelihood estimator $\widehat{\mathbf{V}}_M := \widehat{\Sigma}_{(M)} / (n - |M|)$, with $\widehat{\Sigma}_{(M)} := \mathbf{Y} (\mathbf{I}_n - \mathbf{H}_{(M)}) \mathbf{Y}^\top$ and $\mathbf{H}_{(M)} := \mathbf{X}_M^\top (\mathbf{X}_M \mathbf{X}_M^\top)^{-1} \mathbf{X}_M$. Note that the matrix $\mathbf{H}_{(M)}$ is the orthogonal projection onto the row space of \mathbf{X}_M . Moreover, given the true data generating model (3), it follows that $\mathbb{E}(\widehat{\mathbf{Y}}) = \mathbb{E}(\widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_{M\cdot}) = \mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o} \mathbf{H}_{(M)}$, and so if $M \supseteq M_o$ then $\mathbb{E}(\widehat{\mathbf{Y}}) = \mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o}$. This means that any collection of predictors with linear span containing the oracle predictors, M_o , is as good at explaining variation in the data as the true model (in terms of residual sum of squares). However, such a large/redundant set of predictors lacks efficiency in terms of prediction accuracy. Exploiting this idea, the critical definition supporting our methodology is presented next.

Definition 2.1. A $q \times |M|$ regression coefficient matrix $\mathbf{B}_{\cdot M}$ coupled with an index set $M \subseteq \{1, \dots, p\}$ is said to be ϵ -admissible if $h_\epsilon(\mathbf{B}_{\cdot M}) = 1$ where,

$$h_\epsilon(\mathbf{B}_{\cdot M}) = \mathbb{I} \left\{ \frac{1}{2} \left\| \widehat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} \mathbf{X}_{M\cdot} - \mathbf{B}_{\min} \mathbf{X}) \right\|_{\text{F}}^2 \geq \epsilon \right\} \mathbb{I}\{|M| < n - q\}, \quad (5)$$

where \mathbf{B}_{\min} is the solution to the optimization problem,

$$\arg \min_{\mathbf{B} \in \mathbb{R}^{q \times p}} \left\| \widehat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B} \mathbf{X}) \right\|_{\mathbb{F}}^2 \quad \text{subject to} \quad |\{j : \|\mathbf{B}_j\| \neq 0\}| \leq |M| - 1.$$

Definition 2.1 characterizes a notion of redundancy for any set of predictors, indexed by M . The quantity $\|\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B} \mathbf{X}\|_{\mathbb{F}}^2$ captures the difference in prediction of the model M from all models with fewer predictors. Any model M that is not ϵ -admissible is redundant in the sense that there exists a subset of fewer predictors that approximately linearly spans the same subspace. This very notion of redundancy makes the EAS method different from the traditional regularization-based approaches, where redundancy is expressed as a model containing negligible or zero magnitude regression coefficients. Nonetheless, Definition 2.1 encompasses the traditional notion of redundancy because, if one column of $\mathbf{B}_{\cdot M}$ is equal to zero then $h_{\epsilon}(\mathbf{B}_{\cdot M})$ is 0 for all $\epsilon > 0$. Additionally, as a consequence of the rows of $\mathbf{X}_{M \cdot}$ spanning a finite-dimensional vector space, $h_{\epsilon}(\mathbf{B}_{\cdot M})$ assigns value zero to all models M with $|M| > n$, by definition. As a consequence, the EAS procedure inherently reduces the difficulty of the model selection problem from 2^p candidate models to 2^n , a fact that is fundamental to the scalability of the EAS procedure for high-dimensional settings. Furthermore, if $\mathbf{X}_{M \cdot}$ does not have full row rank, then $h_{\epsilon}(\mathbf{B}_{\cdot M})$ is again zero by its construction.

The above definition of the h -function is well-defined in the sense that for M with $|M| < n - q$, the matrix $\widehat{\Sigma}_{(M)}$ is invertible with probability 1. Lemma 4.2 in Section 4 justifies that the minimum eigenvalue of $\widehat{\Sigma}_{(M)}$ diverges away from 0 for large n . Next, for identifiability of a sparse true model, M_o , the choice of ϵ must not be so large that it classifies M_o as redundant. Conversely, if ϵ is chosen too small, then many redundant models might also satisfy the ϵ -admissibility criterion. It will be seen throughout the remainder of the paper that this trade-off analysis is the crux of the theoretical underpinnings of the EAS approach.

A distinction between the definition of the h -function defined in (5) and the one defined in [53] is the introduction of the empirical error covariance matrix $\widehat{\Sigma}_{(M)}$. Normalization by the square-root of error covariance matrix is common in LASSO-type model selection strategies such as concomitant multi-task regression [33], multivariate square-root LASSO [47, 5] to name a few. For any fixed model M , a common assumption is that the quantity $\|\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B} \mathbf{X}\|_{\mathbb{F}}^2$ is on the order of n (for large n). As a result, the optimal choice of ϵ as derived in [53] turned out to be a function of both n and $|M_o|$. Moreover, the form of their suggested ϵ is somewhat unintuitive as it was derived purely from a theoretical result. On the other hand, since $\widehat{\Sigma}_{(M)}$ is also on the order of n (for large n), adjusting for the inverse square-root of $\widehat{\Sigma}_{(M)}$, as in (5), proportionately scales $\|\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B} \mathbf{X}\|_{\mathbb{F}}^2$. This enables us to choose the threshold ϵ via simple grid search based on some metric such as cross-validation (CV) technique or information criterion (IC), independently of n and $|M|$. We implement a computationally efficient version of the EAS algorithm that bypasses the original repeated sampling based pseudo-marginal Markov chain Monte Carlo (MCMC)

algorithm. A detailed description of the computational procedure we propose is presented in Section 3.

With our ϵ -admissibility notion of redundancy now defined, we are ready to build the statistical framework that will facilitate its use in model selection. The GF inference approach adopted in [53] remains an advantageous pathway for constructing a posterior-like probability distribution over the class of candidate models; one that is principled in the sense of Bernstein-von Mises asymptotics, but also avoids the problem of prior choice/specification.

To introduce the mechanics of GF inference, assume a random variable Z has a forward data generation equation that can be expressed as $Z = G(U, \theta)$, where G is some known deterministic function, U is a pivotal quantity whose distribution is known, and θ is some fixed but unknown parameter(s) of interest in the space Θ . Given an observed data set $\mathbf{z} = (z_1, \dots, z_n)^\top$ of independent instances of the random variable Z , GF inference aims to find the best fitting ϑ such that $\|\mathbf{z} - G(\mathbf{U}, \vartheta)\|$ is minimized, resulting in a random variable $\theta^* := \theta^*(\mathbf{U})$. The distribution of the random variable $\theta^*(\mathbf{U})$ is termed the GF distribution of the unknown parameter θ . When Z is a continuous random variable, under certain regularity conditions [as stated in 24, mostly dealing with the smoothness of G] the GF distribution of θ can be expressed as

$$r(\vartheta \mid \mathbf{z}) = \frac{f(\mathbf{z}, \vartheta)J(\mathbf{z}, \vartheta)}{\int_{\Theta} f(\mathbf{z}, \tilde{\vartheta})J(\mathbf{z}, \tilde{\vartheta}) d\tilde{\vartheta}}, \quad (6)$$

where $f(\mathbf{z}, \vartheta)$ is the likelihood function and

$$J(\mathbf{z}, \vartheta) := D\left(\frac{d}{d\vartheta}G(\mathbf{u}, \vartheta)\Big|_{\mathbf{u}=G^{-1}(\mathbf{z}, \vartheta)}\right), \quad (7)$$

with $D(\mathbf{C}) := \sqrt{\det \mathbf{C}^\top \mathbf{C}}$ for a matrix argument \mathbf{C} . The Jacobian-like quantity $J(\mathbf{z}, \vartheta)$ results from inverting the data generating equation assuming the inverse $G^{-1}(\mathbf{z}, \vartheta)$ exists.

To illustrate, if $Z_1, \dots, Z_n \stackrel{iid}{\sim} N(\mu, \sigma^2)$, then $G(U, (\mu, \sigma^2)) = \mu + \sqrt{\sigma^2}U$ where $U \sim N(0, 1)$. In this case,

$$J(\mathbf{z}, (\mu, \sigma^2)) = D([\mathbf{1}_n - \sigma^{-2}(\mathbf{z} - \mu\mathbf{1}_n)/2]) = \sqrt{n}\sigma^{-2}s_z/2,$$

where $\bar{z} := \sum_{i=1}^n z_i$ and $s_z^2 := \sum_{i=1}^n (z_i - \bar{z})^2$. Then the GF distribution takes the form

$$r(\mu, \sigma^2 \mid \mathbf{z}) \propto (\sigma^2)^{-n/2-1}e^{-\sigma^{-2}\sum_1^n (z_i - \mu)^2/2},$$

which implies that the conditional GF distribution of $(\mu \mid \sigma)$ is $N(\bar{z}, \sigma^2/n)$, and the marginal GF distribution of σ^2 is Inverse-Gamma $\{(n-1)/2, s_z^2/2\}$. This solution is consistent with the posterior distribution of (μ, σ^2) constructed via flat prior specification [page 65, 21].

In contrast to the univariate regression model (as in [53]), however, construction of the GF distribution in the MLR setting is *not* a simple extension and is accompanied by unique challenges, especially so for dealing with the arbitrary covariance matrix of the multivariate response vectors.

The difficulty in deriving and studying an expression for a GF distribution, for most continuous data models, is that it requires deriving a complicated function of partial derivatives with respect to the unknown model parameters. Interesting data models for which the GF density can be expressed analytically (up to a normalizing constant), such as we will show for the MLR with arbitrary coefficient and covariance matrices, are interesting in their own right for their contribution to the growing literature on GF inference. The assumptions and materials for deriving/computing a GF distribution are provided in [24]. In the remainder of this section we provide the details relevant to our methodology.

Given an index set, M , the unknown parameters in model (4) are $\mathbf{B}_{\cdot M}$ and $\mathbf{A}_{(M)}$. As in the GF inference setup, re-express the data generating equation (4) as,

$$\mathbf{Y} = \mathbf{B}_{\cdot M} \mathbf{X}_M + \mathbf{A}_{(M)} \mathbf{U} =: G(\mathbf{U}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})), \quad (8)$$

where $\mathbf{U} \sim \text{Matrix-Normal}_{q,n}(\mathbf{0}, \mathbf{I}_q, \mathbf{I}_n)$. As prescribed in Theorem 1 of [24], the GF density of the parameters $(\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})$ can be expressed as,

$$r(\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)} | \mathbf{Y}) := \frac{f(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) J(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)}))}{\int f(\mathbf{Y}, (\tilde{\mathbf{B}}_{\cdot M}, \tilde{\mathbf{A}}_{(M)})) J(\mathbf{Y}, (\tilde{\mathbf{B}}_{\cdot M}, \tilde{\mathbf{A}}_{(M)})) d(\tilde{\mathbf{B}}_{\cdot M}, \tilde{\mathbf{A}}_{(M)})},$$

where f is the matrix normal likelihood function, and the *Jacobian* term defined in (7). After some routine matrix calculations the Jacobian term reduces to,

$$J(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) = \left(\det \mathbf{A}_{(M)} \mathbf{A}_{(M)}^\top \right)^{-q/2} \left(\det \mathbf{X}_M \mathbf{X}_M^\top \right)^{q/2} \left(\det \hat{\Sigma}_{(M)} \right)^{q/2}.$$

Accordingly, restricting M to the class of ϵ -admissible models yields the GF density,

$$r_\epsilon(\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)} | \mathbf{Y}) \propto e^{-\frac{1}{2} \text{tr}[\mathbf{R}_{(M)} \mathbf{V}_{(M)}^{-1}]} \left(\det \mathbf{X}_M \mathbf{X}_M^\top \right)^{q/2} \left(\det \hat{\Sigma}_{(M)} \right)^{q/2} \times \\ \left(\det \mathbf{V}_{(M)} \right)^{-(n+q)/2} h_\epsilon(\mathbf{B}_{\cdot M}),$$

where $\mathbf{R}_{(M)} := (\mathbf{Y} - \mathbf{B}_{\cdot M} \mathbf{X}_M) (\mathbf{Y} - \mathbf{B}_{\cdot M} \mathbf{X}_M)^\top$. Note the dependence of $r_\epsilon(\cdot | \mathbf{Y})$ on the choice of ϵ . Further, $r_\epsilon(\cdot | \mathbf{Y})$ should not be confused with the notation for a conditional probability density function, but should be understood to reflect the fact that the GF distribution is a function of the observed data \mathbf{Y} .

Moving along, analogous to a Bayesian model selection approach, we construct a probability distribution over all ϵ -admissible index sets as the marginal distribution,

$$r_\epsilon(M | \mathbf{Y}) \propto \int_{\mathbb{R}^{q \times |M|}} \int_{\mathbb{R}^{q \times q}} r_\epsilon(\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)} | \mathbf{Y}) d\mathbf{A}_{(M)} d\mathbf{B}_{\cdot M} \\ = \left(\det \mathbf{X}_M \mathbf{X}_M^\top \right)^{\frac{q}{2}} \left(\det \hat{\Sigma}_{(M)} \right)^{\frac{q}{2}} \int h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{B}_{\cdot M} \times \\ \int \frac{e^{-\frac{1}{2} \text{tr}[\mathbf{R}_{(M)} \mathbf{V}_{(M)}^{-1}]} }{\left(\det \mathbf{V}_{(M)} \right)^{(n+q)/2}} d\mathbf{A}_{(M)}$$

We simplify this expression as equation (9), stated next, and provide a detailed account of the intermediate steps in Section A of the Appendix. The derivation of GF model involve integration over the domain of positive definite matrices and non-trivial matrix algebra that are far more complex than the case of univariate linear regression setting.

$$r_\epsilon(M | \mathbf{Y}) \propto \Gamma_q \left(\frac{n - |M|}{2} \right) \pi^{\frac{q|M|}{2}} \left(\det \widehat{\Sigma}_{(M)} \right)^{-\left(\frac{n - |M| - q}{2} \right)} \mathbb{E} [h_\epsilon(\mathbf{B}_{\cdot M})], \quad (9)$$

where the expectation is taken with respect to the density of matrix t -distribution, i.e.,

$$\mathbf{B}_{\cdot M} \sim T_{q, |M|} \left(n - |M| - q + 1, \widehat{\mathbf{B}}_{\cdot M}, \widehat{\Sigma}_{(M)}, (\mathbf{X}_M \cdot \mathbf{X}_M^\top)^{-1} \right). \quad (10)$$

Note that the GF distribution of $\mathbf{B}_{\cdot M}$ is concentrated around the least squared estimator, $\widehat{\mathbf{B}}_{\cdot M}$, defined previously.

Observe in (9) that for models with $|M| > n - q$, the h -function is zero by definition, and thus $r_\epsilon(M | \mathbf{Y})$ is trivially zero. This probability mass function has the interpretation as the relative likelihood of the model M versus that of all other candidate models in the class of ϵ -admissible models. It becomes clear from the expression (9) that $r_\epsilon(M | \mathbf{Y})$ is largely driven by the inverse of the empirical error covariance matrix, and that the h -function delivers a multiplicative effect on the probability. For a large redundant model, M , we expect that the determinant of the empirical error covariance is small relative to that of M_o , and so we leverage the choice of ϵ such that $\mathbb{E} [h_\epsilon(\mathbf{B}_{\cdot M})]$ controls the value of $r_\epsilon(M | \mathbf{Y})$. This insight is formalized in Section 4.

As we illustrate in the remainder of this paper, the GF mass function $r_\epsilon(M | \mathbf{Y})$ serves as a vehicle for model selection and inference. In the next section we discuss the details of the computations, and provide an algorithm to generate samples from this GF distribution.

3. Model estimation and computational techniques

In order to generate samples from $r_\epsilon(M | \mathbf{Y})$ we must be able to compute $\mathbb{E} [h_\epsilon(\mathbf{B}_{\cdot M})]$. Although the expectation is with respect to a matrix t -distribution, the complex expression for $h_\epsilon(\mathbf{B}_{\cdot M})$ makes the form of its expectation intractable, and so standard MCMC techniques do not apply. This issue is typical of all the previously developed EAS implementations. The MLR EAS analogue of the previous EAS approaches is to employ a pseudo-marginal MCMC algorithm by estimating $\mathbb{E} [h_\epsilon(\mathbf{B}_{\cdot M})]$ with the average of a large number of random samples from the GF distribution of $\mathbf{B}_{\cdot M}$ (i.e., its matrix t -distribution). An important remark is that in contrast to the previous EAS articles, in our empirical investigations we find that rather than generating a sample of matrices $\mathbf{B}_{\cdot M}$ as in (10) for approximating $h_\epsilon(\mathbf{B}_{\cdot M})$ with an a sample mean, it suffices to take $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$ as a point estimate of $\mathbb{E} [h_\epsilon(\mathbf{B}_{\cdot M})]$, where $\widehat{\mathbf{B}}_{\cdot M}$ is the least squares estimator for model M . This is likely partly due to our construction for the h function having a better scaling with ϵ than in earlier developments

of EAS approaches, and the fact that the distribution in (10) is centered at $\widehat{\mathbf{B}}_{.M}$. This adjustment makes the implementation of EAS highly efficient, and is supported by the competitive performance exhibited in extensive numerical studies, summarized in Section 5. We briefly discuss the algorithms of EAS for MLR case, next.

From Definition 2.1, evaluating $h_\epsilon(\widehat{\mathbf{B}}_{.M})$ can be formulated as the mixed integer quadratic program (MIQP) with quadratic constraints,

$$\text{vec}(\mathbf{B}_{\min}) = \arg \min_{\mathbf{b}, z_1, \dots, z_p} \left\{ \frac{1}{2} \mathbf{b}^\top \mathbf{Q}_{(M)} \mathbf{b} - \mathbf{a}_{(M)}^\top \mathbf{b} \right\},$$

subject to $\mathbf{b} \in \mathbb{R}^{qp}$, $z_j \in \{0, 1\}$ for $j \in \{1, \dots, p\}$,

$$\mathbf{b}^\top \mathbf{C}_j \mathbf{b} \leq \mathcal{M}_U z_j^2, \quad \text{and} \quad \sum_{j=1}^p z_j \leq |M| - 1,$$

where $\mathbf{Q}_{(M)} := (\mathbf{X}\mathbf{X}^\top) \otimes \widehat{\Sigma}_{(M)}^{-1}$, $\mathbf{a}_{(M)} := (\mathbf{X}\mathbf{X}_M^\top \cdot \otimes \widehat{\Sigma}_{(M)}^{-1}) \text{vec}(\widehat{\mathbf{B}}_{.M})$, \mathbf{C}_j is a block diagonal matrix with p blocks of $q \times q$ zero matrices, except \mathbf{I}_q in the j -th block, and $\mathcal{M}_U > 0$ is a properly chosen constant. In particular, the quantity \mathcal{M}_U must be chosen large enough so that $\mathcal{M}_U > \max_j \|(\mathbf{B}_{\min})_j\|_2^2$ [6]. Further, since \mathbf{B}_{\min} is not a-priori known, [6] provides a data-driven formula to specify \mathcal{M}_U in the MIQP, which can be solved with any MIQP solver, such as CPLEX.

Input: Input a model with index set M , least squared estimator $\widehat{\mathbf{B}}_{.M}$, the full design matrix \mathbf{X} , $\widehat{\Sigma}_{(M)}$, a pre-specified $\epsilon > 0$ and an initial solution \mathbf{B}_{init} with number of columns with non-zero norm less than $|M| - 1$;

Output: Value of $h_\epsilon(\widehat{\mathbf{B}}_{.M})$;

- 1 Calculate $L = \lambda_{\max}(\mathbf{X}\mathbf{X}^\top) \lambda_{\min}^{-1}(\widehat{\Sigma}_{(M)})$ and set $\mathbf{B}_{cur} = \mathbf{B}_{init}$;
- 2 Calculate the objective function

$$g(\mathbf{B}_{cur}) = \frac{1}{2} \|\widehat{\Sigma}_{(M)}^{-1/2} (\widehat{\mathbf{B}}_{.M} \mathbf{X}_M \cdot - \mathbf{B}_{cur} \mathbf{X})\|_F^2;$$
- 3 **while** *diff* > *threshold* or $g(\mathbf{B}_{cur}) > \epsilon$ **do**
- 4 Calculate $\mathbf{B} = \mathbf{B}_{cur} - \frac{1}{L} \widehat{\Sigma}_{(M)}^{-1} (\mathbf{B}_{cur} \mathbf{X}\mathbf{X}^\top - \widehat{\mathbf{B}}_{.M} \mathbf{X}_M \cdot \mathbf{X}^\top)$;
- 5 Obtain indices $i_1, i_2, \dots, i_{|M|-1}$ such that

$$\|\mathbf{B}_{i_1}\| \geq \|\mathbf{B}_{i_2}\| \cdots \geq \|\mathbf{B}_{i_{|M|-1}}\| \geq \|\mathbf{B}_{i_{|M|}}\| \cdots \geq \|\mathbf{B}_{i_p}\| ;$$
- 6 Set $\mathbf{B}_j = 0$ for all $j \in \{1, 2, \dots, p\} \setminus \{i_1, i_2, \dots, i_{|M|-1}\}$;
- 7 Calculate *diff* = $|g(\mathbf{B}_{cur}) - g(\mathbf{B})|$;
- 8 Update $\mathbf{B}_{cur} = \mathbf{B}$;
- 9 **end**
- 10 **return** $h_\epsilon(\widehat{\mathbf{B}}_{.M}) = \mathbf{I}(g(\mathbf{B}_{cur}) > \epsilon)$;

Algorithm 3.1: Pseudocode for projected gradient descent to compute $h_\epsilon(\widehat{\mathbf{B}}_{.M})$.

Although a single MIQP is typically fast to solve in practice, we need to compute $h_\epsilon(\widehat{\mathbf{B}}_{.M})$ for different models M at each step of the MCMC, so fur-

ther streamlining of the computations are needed. First, observe that a solution, \mathbf{B}_{\min} , to the MIQP is not always necessary to evaluate $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$; rather $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M}) = 0$ if there exists any \mathbf{B} (satisfying the MIQP constraints) such that $\frac{1}{2} \|\widehat{\Sigma}_{(M)}^{-1/2}(\widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_M - \mathbf{B} \mathbf{X})\|_F^2 < \epsilon$. If this ‘stopping’ condition is met prior to obtaining \mathbf{B}_{\min} , then the MIQP solver can be terminated early. Second, as an alternative to an explicit MIQP solver, a discrete first-order gradient-descent based algorithm proposed in [6] can be implemented for a crude but super-efficient computation of $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$, when $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M}) = 0$. This ‘projected gradient-descent’ algorithm is advocated as a warm start to the MIQP in [6], and the pseudocode for our implementation of it is given in Algorithm 3.1. Note that the gradient of objective function in the optimization problem in Definition 2.1 is Lipschitz continuous with Lipschitz constant $L = \|\widehat{\Sigma}_{(M)}^{-1}\| \|\mathbf{X} \mathbf{X}^\top\|$. For $\widehat{\mathbf{B}}_{\cdot M}$ that are *not* ϵ -admissible, we observed in empirical experimentation that by initializing Algorithm 3.1 at $\widehat{\mathbf{B}}_{\cdot M}$, with its column having minimum norm set to zero, it usually finds a solution to determine that $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$ is zero within a few iterations.

Input: An index set M and $\epsilon > 0$
Output: A new index set M^*

- 1 Calculate

$$\widetilde{M} = \begin{cases} M \cup \{\text{a new covariate}\} & w.p. \ 1/3 \\ M \setminus \{\text{a existing covariate}\} & w.p. \ 1/3; \\ M \setminus \{\text{a existing covariate}\} \cup \{\text{a new covariate}\} & w.p. \ 1/3 \end{cases}$$
- 2 Calculate LS estimator of $\widehat{\mathbf{B}}_{\cdot M}$ and $\widehat{\Sigma}_{(M)}$ corresponding to M ;
- 3 Calculate LS estimator of $\widehat{\mathbf{B}}_{\cdot \widetilde{M}}$ and $\widehat{\Sigma}_{(\widetilde{M})}$ corresponding to the proposal \widetilde{M} ;
- 4 Compute $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$ and $h_\epsilon(\widehat{\mathbf{B}}_{\cdot \widetilde{M}})$ using Algorithm 3.1;
- 5 Calculate $\widehat{r}_\epsilon(M | \mathbf{Y})$ and $\widehat{r}_\epsilon(\widetilde{M} | \mathbf{Y})$ as in (11) ;
- 6 $M^* = \begin{cases} \widetilde{M} & w.p. \ \rho(M, \widetilde{M}) \\ M & w.p. \ 1 - \rho(M, \widetilde{M}) \end{cases}$, where $\rho(M, \widetilde{M}) := \min \left\{ \frac{\widehat{r}_\epsilon(\widetilde{M} | \mathbf{Y})}{\widehat{r}_\epsilon(M | \mathbf{Y})}, 1 \right\}$;

Algorithm 3.2: Pseudocode for one step of the MCMC algorithm to estimate $r_\epsilon(M | \mathbf{Y})$. Note that weights can be used for randomly adding/dropping/switching covariates in the proposed index set \widetilde{M} in line 2 (e.g., correlation-based weights). In that case, the Metropolis-Hasting ratio $\rho(M, \widetilde{M})$ in line 7 needs to be updated accordingly.

Now that we have a computationally efficient algorithm for computing the $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$, the remaining task is to demonstrate the mechanism for generating samples from the GF distribution of M . Estimating $\mathbb{E}[h_\epsilon(\mathbf{B}_{\cdot M})]$ by $h_\epsilon(\widehat{\mathbf{B}}_{\cdot M})$, the GF probability mass function $r_\epsilon(M | \mathbf{Y})$ can be approximated as

$$\widehat{r}_\epsilon(M | \mathbf{Y}) := \Gamma_q \left(\frac{n - |M|}{2} \right) \pi^{\frac{q|M|}{2}} \left(\det \widehat{\Sigma}_{(M)} \right)^{-\left(\frac{n - |M| - q}{2} \right)} h_\epsilon \left(\widehat{\mathbf{B}}_{\cdot M} \right). \quad (11)$$

We present the pseudocode for our implementation in Algorithm 3.2. We demonstrate empirically in Section 5 that this approximation gives results that are competitive with the state-of-the-art Bayesian and frequentist methods for the MLR, both in terms of performance and computation time.

4. Theoretical results

The main objective of this section is to establish the consistency of our model selection procedure, particularly in the high-dimensional setting (i.e., $p \gg n$) with the assumption that the true model is sparse. We begin by stating and describing essential conditions and necessary supporting results to show that $r_\epsilon(M_o | \mathbf{Y})$ converges in probability to 1 as $n \rightarrow \infty$. Our strong model selection consistency result is stated as Theorem 4.10. Throughout this section we a-priori fix the following values. Let $\underline{\lambda}_v$ and $\bar{\lambda}_v$ be the minimum and maximum eigenvalues of the true covariance matrix, $\mathbf{V}_{(M_o)}^0$, respectively. Denote by $\mathbb{P}_y(\cdot)$ be the probability measure associated with the sampling distribution of the response \mathbf{Y} , as in (3), and denote by $\mathbb{P}(\cdot)$ the probability measure associated with the GF distribution of the parameters. Similarly, denote by $\mathbb{E}_y(\cdot)$ the expectation with respect to the sampling distribution of the response \mathbf{Y} , as in (3), and denote by $\mathbb{E}(\cdot)$ the expectation with respect to the GF distribution of the parameters.

The major theoretical intricacies that we deal with while extending from the high-dimensional univariate linear regression case are, first, the residual sum of squares in the multivariate linear model is no longer a scalar, indeed a $q \times q$ matrix. As evident from the equation (9), in order to establish asymptotic variable selection consistency of the true model, we must derive the concentration bound on the ratio of the determinant of residual sum of squares of matrix for any arbitrary model to that for the true model. This calls for a lower bound on the size of minimum eigenvalue of the $\hat{\Sigma}_{(M)}$, for which the existing theoretical results are sparse, compared to the well-established standard chi-squared tail bounds that apply to the RSS in univariate setting. Second, the non-asymptotic bound for the ratio of the determinant of error covariance matrix is sharper in the sense that it does not pivot upon the growth of $\Delta_{(M)} := \mathbf{B}_{M_o}^0 \mathbf{X}_{M_o} (\mathbf{I}_n - \mathbf{H}_{(M)}) \mathbf{X}_{M_o}^\top \mathbf{B}_{M_o}^{0\top}$. Consequently, we establish strong variable selection consistency of the true model without imposing any rate on $\Delta_{(M)}$ [53, first part of Condition 3.2], quite distinctly from the univariate paper. Third, the concentration bounds for the $\mathbb{E}(h_\epsilon(\mathbf{B}_{M_o}))$ are derived explicitly as a function of n , in contrast to Theorem 3.7 and 3.8 of [53] where the bounds are expressed as a function of ϵ , which demands for an additional non-intuitive assumption [53, Condition 3.4] on the rate of ϵ . Theorem 4.10 in our paper guarantees that such a condition is not essential to establish variable selection consistency as long as ϵ satisfies Condition 4.6 and 4.8, which are fundamental to the EAS methodology.

Condition 4.1 requires that the true covariance matrix of the response is positive-definite and finite, which implies that none of the components of the multivariate response are degenerate and they all have finite second moments.

Since the dimension of the multivariate response, q , is fixed, this assumption is rather routine.

Condition 4.1 (*Non-singularity of true covariance*). *The dimension of the multivariate response vector, q , is fixed, and $0 < \underline{\lambda}_v < \bar{\lambda}_v < \infty$.*

In our methods, it is important that $\mathbf{V}_{(M_o)}^0$ is positive-definite so that there exists a positive-definite, consistent estimator of it, for example, $\widehat{\Sigma}_{(M_o)}/(n - |M_o|)$. This estimator plays an essential role in our definition of the h function (among other roles). In particular, Condition 4.1 makes it possible that, for large n , the minimum eigenvalue of $\widehat{\Sigma}_{(M)}/(n - |M|)$ is bounded away from 0 with high probability, for an important class of models. This fact is established in Lemma 4.2, presented next.

Lemma 4.2. *Assume the data generating model (3), and that Condition 4.1 holds. Then for sufficiently large n , and for any fixed $0 < \tau < 1$, and for any $M \subseteq \{1, 2, \dots, p\}$ satisfying $n > |M| + 2q$,*

$$\mathbb{P}_y \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) \geq \tau(n - |M|)\underline{\lambda}_v \right) > 1 - e^{-(1-\sqrt{\tau})^2(n-|M|)/2}$$

Condition 4.3 specifies that the number of predictors are allowed to grow at sub-exponential rate with the sample size, n , ensuring that our method is suitable to perform in the high-dimensional setting. Model selection consistency with this size of p relative to n is on par with the state-of-the-art results in the literature [1].

Condition 4.3. *For some fixed $\alpha \in (0, 1)$, $\log p = o(n^{1-\alpha})$.*

For a given model M , if $|M|$ is on the order of n , then the row space of \mathbf{X}_M might span \mathbb{R}^n leading to a rank deficient empirical error covariance matrix, $\widehat{\Sigma}_{(M)}$. Accordingly, $r_\epsilon(M | \mathbf{Y})$ from equation (9) will be undefined in this case. Since we assume that the data arise from the non-degenerate statistical model (3), we must exclude index sets M with $n^\alpha < |M| < n - q$ for some fixed $\alpha \in (0, 1)$ arising in Condition 4.3. Recall, however, that the h -function already assigns the value 0, by definition, to M with $|M| \geq n - q$, and that q is small and fixed. Throughout the remainder of this section, we assume $\alpha \in (0, 1)$ to be some a-priori fixed value with $|M_o| \leq n^\alpha$. Condition 4.3 is quintessential for LASSO to achieve variable selection consistency in high-dimensional settings [Theorem 3 of 56]. It will be observed in the coming results that the fraction α can be interpreted as a tuning parameter that balances the maximum model size to be considered versus the rate of convergence of the EAS procedure.

Moving along, Condition 4.4 ensures that the design matrix for the true model has full row rank. This type of restricted eigenvalue condition is routinely needed in the variable selection literature; e.g., see Condition 6 of [56] and [26] in the context of LASSO.

Condition 4.4. *For the true model M_o , $\mathbf{X}_{M_o} \cdot \mathbf{X}_{M_o}^\top$ is non-singular.*

Next, in order to show that $r_\epsilon(M_o | \mathbf{Y}) \rightarrow 1$ in probability, we must show that $r_\epsilon(M | \mathbf{Y})/r_\epsilon(M_o | \mathbf{Y}) \rightarrow 0$ in probability, at a rate vanishing faster than

2^{-n^α} uniformly for every model $M \neq M_o$ with $|M| \leq n^\alpha$. Recall that the probability mass function $r_\epsilon(M \mid \mathbf{Y})$, in equation (9), is proportional to a polynomial of the inverse of the determinant of the empirical error covariance matrix. That being so, we must bound the ratio of determinant, as in Theorem 4.5. This ratio is analogous to the ratio of RSS that commonly appears in univariate model selection problems, though, the multivariate situation is much more complicated requiring delicate handling of minimum eigenvalues close to zero. That being so, this result is interesting in its own right for (high-dimensional) MLR.

For any model M with $|M| \leq n^\alpha$, it is understood that the determinant of $\widehat{\Sigma}_{(M_o)}$ as a proportion of that of $\widehat{\Sigma}_{(M)}$ will behave differently depending on whether $M \subsetneq M_o$ or $M \not\subsetneq M_o$. In the first case, when $M \subsetneq M_o$, the ratio will be strictly less than 1 since M is missing at least one oracle predictor. Conversely, the extreme scenario in the other case is that $M \supset M_o$, in which case the ratio exceeds 1, but by some bound that converges to 1 as n tends to infinity.

Theorem 4.5. *Assume Conditions 4.1, 4.3, and 4.4. Then for sufficiently large n , the following approximations hold.*

Case 1: This case pertains to the models $M \subsetneq M_o$.

$$\mathbb{P}_y \left(\bigcap_{M: M \subsetneq M_o} \left\{ \mathbf{Y} : \left(\frac{\det \widehat{\Sigma}_{M_o}}{\det \widehat{\Sigma}_M} \right)^{\frac{n-|M|-q}{2}} \leq e^{-qn^\alpha \log |M_o|} \right\} \right) \geq 1 - V_{1,n},$$

where,

$$V_{1,n} := |M_o| q \exp \left(-\frac{\xi_{n,|M_o|} n^\alpha \log |M_o| \lambda_v}{2\lambda_v} - 0.09(n - |M_o|) + |M_o| \log(|M_o|) \right) + 2|M_o| \exp(-0.04(n - |M_o|) + |M_o| \log |M_o|),$$

with $\xi_{n,|M|} := 1 - \frac{2n^\alpha \log |M_o|}{n - |M| - q}$ such that $\xi_{n,|M|} \in (0, 1)$ for large n .

Case 2: This case pertains to the models $M \not\subsetneq M_o$ such that $|M| \leq n^\alpha$.

$$\mathbb{P}_y \left(\bigcap_{\substack{M: M \not\subsetneq M_o \\ |M| \leq n^\alpha}} \left\{ \mathbf{Y} : \left(\frac{\det \widehat{\Sigma}_{M_o}}{\det \widehat{\Sigma}_M} \right)^{\frac{n-|M|-q}{2}} \leq e^{q(n^\alpha \log(n - |M_o|) + |M| \log p)} \right\} \right) \geq 1 - V_{2,n},$$

where,

$$V_{2,n} := 2q \exp \left(-\frac{n^\alpha}{2} \log(n - |M_o|) + \alpha \log n + \frac{(|M_o| + 1)}{2} \log(\zeta_{n,n^\alpha}) \right),$$

and $\zeta_{n,|M|} := 1 + \frac{2(n^\alpha \log(n - |M_o|) + |M| \log p)}{n - |M| - q}$.

Observe that both quantities $V_{1,n}$ and $V_{2,n}$ vanish exponentially fast for large n , by Condition 4.3. There are two key facts that we learn from Theorem 4.5. The first is that the ratio of the determinants, of the empirical error covariances raised to the power on the order of n , will drive $r_\epsilon(M \mid \mathbf{Y})/r_\epsilon(M_o \mid \mathbf{Y})$ to 0 for $M \subsetneq M_o$ (i.e., Case 1). The second fact is that the ratio of the determi-

nants, of the empirical error covariances raised to the power on the order of n , will perhaps grow at a sub-exponential rate for large n , for $M \not\subseteq M_o$ such that $|M| \leq n^\alpha$ (i.e., Case 2). As such, the role of the h -function is to control the explosive nature of these ratios for models with redundant predictors. The following two theorems establish that $\mathbb{E}(h_\epsilon(\mathbf{B}_{.M}))$ and $\mathbb{E}(h_\epsilon(\mathbf{B}_{.M_o}))$ are adept at accomplishing this task. Sufficient conditions on the choice of ϵ are stated in the order that they are needed.

Condition 4.6 (ϵ -admissibility). *The size of the true model $|M_o|$ is less than n^α . Moreover, for large n , the true model M_o satisfies,*

$$\frac{1}{36q(n - |M_o|)} \left\| \left(\mathbf{V}_{(M_o)}^0 \right)^{-1/2} \left(\mathbf{B}_{M_o}^0 \mathbf{X}_{M_o} - \tilde{\mathbf{B}}_{\min} \mathbf{X} \right) \right\|_F^2 > \epsilon,$$

where $\tilde{\mathbf{B}}_{\min}$ is the solution to the optimization problem,

$$\tilde{\mathbf{B}}_{\min} := \arg \min_{\mathbf{B} \in \mathbb{R}^{q \times p}} \left\| \left(\mathbf{V}_{(M_o)}^0 \right)^{-1/2} \left(\mathbf{B}_{M_o}^0 \mathbf{X}_{M_o} - \mathbf{B} \mathbf{X} \right) \right\|_F^2,$$

subject to $|\{j : \|\mathbf{B}_j\| \neq 0\}| \leq |M_o| - 1$.

Condition 4.6 provides the maximum rate of growth for the size of the true model. This is analogous to the *sparsity* assumption for the LASSO [Condition 7 of 56]. It also furnishes an upper bound for the choice of ϵ that is sufficient for the identifiability of the true model and coefficients, as in Definition 2.1 of the h function. Notice that given an ϵ , the smaller the norm of the regression coefficient matrix $\mathbf{B}_{M_o}^0$, the more difficult it becomes to identify the true model as ϵ -admissible. It is in this sense that ϵ -admissibility defines redundancy both in the sense of correlated predictors and in the sense of predictors with weak signal (after scaling for the response covariance). This is related to the ‘beta-min’ condition discussed for variable selection via LASSO ([Section 7.4 of 11] and [Condition 8 of 56]).

With the addition of Condition 4.6, Theorem 4.7 ensures that the oracle model is ϵ -admissible. In our proof strategy of Theorem 4.10 (Appendix B), this theorem provides a non-asymptotic probabilistic guarantee that $\mathbb{E}(h_\epsilon(\mathbf{B}_{.M_o}))$ in the denominator of $r_\epsilon(M | \mathbf{Y})/r_\epsilon(M_o | \mathbf{Y})$ is bounded away from zero, so long as ϵ is not too large.

Theorem 4.7. *Assume Conditions 4.1 and 4.4. Then, for every $\epsilon > 0$ satisfying Condition 4.6,*

$$\begin{aligned} \mathbb{P}_y \left[\mathbb{E}(h_\epsilon(\mathbf{B}_{.M_o})) \geq 1 - \exp \left(-\frac{\epsilon(n - |M_o|)}{36} + \frac{q|M_o|}{2} \right) - 2e^{-\frac{1}{8}\{\sqrt{n-|M_o|} - 2\sqrt{q}\}^2} \right] \\ \geq 1 - V_{3,n}, \end{aligned}$$

where,

$$\begin{aligned} V_{3,n} := \exp \left(-\frac{\epsilon(n - |M_o|) \lambda_v}{4 \bar{\lambda}_v} + \frac{q|M_o|}{2} \right) \\ + \exp(-0.04(n - |M_o|)) + \exp(-0.15q(n - |M_o|)). \end{aligned}$$

Lastly, to justify that $r_\epsilon(M | \mathbf{Y})/r_\epsilon(M_o | \mathbf{Y}) \rightarrow 0$ for all redundant models, it remains to establish that $\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M))$ vanishes rapidly for all models $M \not\subseteq M_o$ with $|M| \leq n^\alpha$ (recall the cases in Theorem 4.5). This brings us to the final major supporting result, Theorem 4.9, for establishing strong model selection consistency. However, in contrast to the the upper bound condition on ϵ in Condition 4.6, a lower bound condition on ϵ is sufficient to ensure that $h_\epsilon(\mathbf{B}_\cdot M)$ assigns negligible probability mass to redundant models via $r_\epsilon(M | \mathbf{Y})$.

Condition 4.8 (Redundancy). For any model M with $M \not\subseteq M_o$ with $|M| \leq n^\alpha$, for large n ,

$$\frac{9}{\lambda_v(n - |M|)} \|\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)})\|_{\mathbb{F}}^2 < \epsilon,$$

where $\mathbf{H}_{(M)(-1)} := \mathbf{H}_{(M \setminus \{j^*\})}$ is the projection matrix for the size that is constructed after omitting the predictor j^* from the model M that minimizes

$$j^* = \arg \min_{j \in M} \|\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M \setminus \{j\})})\|_{\mathbb{F}}^2.$$

Condition 4.8 is sufficient for showing that $\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M)) \rightarrow 0$ in probability for all redundant models, and further characterizes the non- ϵ -admissible notion for redundancy. The quantity on the left side of the condition is the mean difference in the prediction between models M and $M \setminus \{j^*\}$. Condition 4.8 implies that models M with $M \not\subseteq M_o$ are redundant in the sense that they contain at least one predictor whose omission will not change the mean predicted response by more than ϵ , as measured by the properly scaled squared Frobenius norm. This requires that none of the predictors in model $M \setminus M_o$ can be replaced by some predictors in the true model to provide a significantly better prediction than ϵ (in the appropriate scale). Intuitively, this means that the correlations between the predictors in the true model and the ones not in the true model cannot be large. Resembling the notion of the *irrepresentability* condition necessary for LASSO model selection consistency [56], this is to say that the irrelevant covariates cannot be well-represented by any of the covariates in the true model.

Condition 4.6 coupled with Condition 4.8 provides the crucial interval for the choice of ϵ within which the oracle model is identifiable and the EAS procedure achieves strong model selection consistency. Notably, due to the appropriate scaling of quantities in the h function, this interval neither depends on the sample size n nor the size of the model M .

Theorem 4.9. Assume Conditions 4.1, 4.3, and 4.4. Then, for sufficiently large n , and for every $\epsilon > 0$ satisfying Condition 4.8,

$$\mathbb{P}_y \left[\bigcap_{\substack{M \not\subseteq M_o \\ |M| \leq n^\alpha}} \left\{ \mathbf{Y} : \mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M)) \leq \exp\left(-\frac{\epsilon(n - |M|)}{36} + \frac{q|M|}{2}\right) + 2e^{-\frac{1}{8}\{\sqrt{n-|M|}-2\sqrt{q}\}^2} \right\} \right] \geq 1 - V_{4,n},$$

where,

$$V_{4,n} := n^\alpha \left\{ \exp \left(-\frac{\epsilon(n - n^\alpha) \lambda_v}{36 \bar{\lambda}_v} + \frac{qn^\alpha}{2} + n^\alpha \log p \right) + 2 \exp(-0.04(n - n^\alpha) + n^\alpha \log p) \right\}.$$

Theorem 4.9 is a non-asymptotic concentration bound for $\mathbb{E}(h_\epsilon(\mathbf{B}_M))$ that applies uniformly over all model $M \not\subseteq M_o$. This is the critical theoretical aspect of the h function that compensates for the explosive nature of the ratios of the determinants of the empirical error covariances raised to the power on the order of n , uniformly over all models $M \not\subseteq M_o$ such that $|M| \leq n^\alpha$, as exhibited in Case 2 of Theorem 4.5.

To this point in the article, sufficient analysis has been constructed to argue the *pairwise* model selection consistency result that $r_\epsilon(M | \mathbf{Y})/r_\epsilon(M_o | \mathbf{Y}) \rightarrow 0$ in probability for any $M \neq |M_o|$ with $|M| \leq n^\alpha$. For the case when p is fixed, this also implies *strong* model selection consistency. In the case when $p \rightarrow \infty$ and particularly for $p \gg n$, however, further justification is required because the number of candidate models to consider is 2^{n^α} . Theorems 4.7 and 4.9 are able to manage this exponential-sized class of candidate model with the essential attribute that they provide concentration inequalities of tails that are uniform and vanish exponentially fast in n . This fact is stated as our main result, Theorem 4.10.

Theorem 4.10. *Assume the data generating model (3), and suppose that Conditions 4.1, 4.3 and 4.4 are satisfied. Then for every $\epsilon > 0$ satisfying Conditions 4.6 and 4.8,*

$$\frac{r_\epsilon(M_o | \mathbf{Y})}{\sum_{M: |M| \leq n^\alpha} r_\epsilon(M | \mathbf{Y})} \xrightarrow{\mathbb{P}_y} 1,$$

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

The proof of Theorem 4.10 and the proofs of all other results are organized in the Supplementary Material. Note that Theorem 4.10 is the only non-asymptotic result in our theoretical developments, and so as long as the conditions are satisfied, it is expected that it is reasonably illustrative of the performance of our constructed EAS procedure on observed data. We provide evidence to substantiate this claim in finite sample numerical studies, presented next in Section 5.

5. Numerical results

In this section we demonstrate the performance of our EAS method in comparison to the state-of-the-art variable selection procedures for MLR. Very recently [1] developed the MBSP method that is equipped to perform variable selection for MLR. They demonstrate a distinctly superior performance of MBSP over all the existing methods, especially in a high-dimensional setting. To make standard the comparison between the MBSP and EAS approaches, we mimic the exact same synthetic data simulation study design constructed in [1].

The simulation design can be broadly categorized into three parts, *low dimensional (LD)* ($n > p$), *high-dimensional (HD)* ($p > n$) and *ultra high-dimensional (UHD)* ($p \gg n$). Two sub-categories are considered within each of these categories, to analyze performance for varying sizes of the true model, $|M_o|$, versus the total number of predictors, p . Within each category the dimension of the multivariate response, q , is also varied to study the effect of q on the model selection performance. In total, there are six experiments, summarized in Table 1.

For each of the first six simulation designs, we generate synthetic data by the following mechanism: The n columns of the design matrix \mathbf{X} are sampled from a multivariate normal distribution with mean zero and covariance matrix $\mathbf{\Gamma}$, that has an AR(1) structure with correlation coefficient 0.5 (i.e., $\Gamma_{ij} = 0.5^{|i-j|}$, for $i, j \in \{1, \dots, p\}$). The true model M_o is constructed by randomly selecting $|M_o|$ elements from $\{1, \dots, p\}$. Once the true model M_o is constructed, each component of the $q \times |M_o|$ true regression coefficient matrix $\mathbf{B}_{M_o}^0$ is set as a value generated from the random variable $U + I(U > -0.5)$, with $U \sim \text{Uniform}(-5, 4)$, so that the values always lie within $[-5, -0.5] \cup [0.5, 5]$. The response vectors $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are independently generated from a multivariate normal distribution with mean $\mathbf{B}_{M_o}^0 \mathbf{X}_{M_o}$ and covariance $\mathbf{V}_{(M_o)}^0$, where $\mathbf{V}_{(M_o)}^0$ also has an AR(1) structure with $\mathbf{V}_{(M_o),ij}^0 = \sigma^2 0.5^{|i-j|}$ for $i, j \in \{1, \dots, q\}$ and $\sigma^2 = 2$.

TABLE 1
n is the sample size, *p* is the number of predictors, *q* is the dimension of multivariate response vectors, and $|M_o|$ is the size of the true model.

Dimension	Sparsity	n	p	q	$ M_o $
LD ($n > p$)	Sparse	60	30	3	5
	Dense	80	60	6	40
HD ($p > n$)	Sparse	50	200	5	20
	Dense	60	100	6	40
UHD ($p \gg n$)	Ultra-sparse	100	500	3	10
	Sparse	150	1000	4	50

In order to further investigate the performance of the EAS method under challenging scenarios, we conduct additional experiments beyond those presented in [1]. Specifically, we consider three additional experimental settings, described in Table 2. In the first experiment, we increase the dimension of the response variable to $q = 60$, which is ten times larger than the dimension considered in Table 1. In the second experiment, we increase the difficulty of the design matrix by introducing a non-decaying structure for $\mathbf{\Gamma}$, while keeping the error covariance $\mathbf{V}_{(M_o)}^0$ fixed at AR(1). In the third experiment, we introduce a dense structure for the error covariance matrix.

The EAS procedure is implemented by computing Algorithm 3.2 described in Section 3 to draw MCMC samples from the space of all 2^p candidate models. Observe that Algorithm 3.2 is developed to work for a fixed ϵ . We propose two methods for selecting the tuning parameter ϵ ; (i) 10-fold CV and (ii) via Bayesian IC (BIC) by searching over a pre-specified grid of ϵ values. For both the CV and BIC routines, we take a uniform grid of 24 possible values for ϵ , from 0.05 to 10 in all six experiments. In the CV procedure, for each of the 10 folds we

TABLE 2

The notation is defined as follows: n is the sample size, p is the number of predictors, q is the dimension of multivariate response vectors, $|M_o|$ is the size of the true model, and Γ_{ij} and $\mathbf{V}_{(M_o),ij}^0$ are (i, j) th element of Γ (i.e., the variance of the columns of the design matrix of \mathbf{X}) and $\mathbf{V}_{(M_o)}^0$ (i.e., the error covariance matrix), respectively.

Dimension	Set up	q	Γ_{ij}	$\mathbf{V}_{(M_o),ij}^0$
$n = 150,$ $p = 1000,$ $M_o = 50$	Large q	60	$0.5^{ i-j }$	$2 \times 0.5^{ i-j }$
	Large q and Non-decaying correlation of design matrix	60	$0.5\{1 + \mathbb{I}(i = j)\}$	$2 \times 0.5^{ i-j }$
	Large q , non-decaying correlation of design matrix, and dense error covariance matrix	60	$0.5\{1 + \mathbb{I}(i = j)\}$	$1 + \mathbb{I}(i = j)$

implement our EAS method on the training set by running the MCMC 500 steps, discarding the first 200 steps, and evaluating the performance on the validation set, as follows. The initial estimates from the multivariate LASSO (MLASSO) [19] serve as the weights for proposing/removing predictors in the MCMC algorithm. The expectation of $h_\epsilon(\mathbf{B}_M)$ for any model M is approximated by evaluating the h at the least square estimator $\widehat{\mathbf{B}}_M$, which makes the computation very fast compared to the previous version of EAS procedures. The MAP estimated model from the MCMC sample is taken as a point estimator to compute the mean squared prediction error (MSPE) on the validation set. The optimal ϵ is chosen as the one that minimizes the average of the MSPE over the 10 folds. Finally, we re-run Algorithm 3.2 on the entire dataset using the optimal selected ϵ for 10,000 MCMC steps and discard the first 5,000.

For the BIC procedure, the computational cost is much less. In this case, for every ϵ in the grid, we run Algorithm 3.2 for 5,000 steps, discard the initial 2,000 in obtaining the MAP estimated model, and compute the BIC for the MAP model. The ϵ corresponding to the minimum BIC value is selected as optimal. The advantage of using BIC is that we do not need to run the algorithm again for the optimally chosen ϵ , we can simply use the MCMC chain from the initial runs as our estimated sample for the chosen ϵ .

We compare the performance of our EAS method with (1) the MBSP approach as implemented in the R package MBSP [2]; (2) the multivariate group lasso with spike and slab prior (MBGL-SS) method as implemented in the R package MBSGS [32] with the natural grouping (i.e., each predictor represents one group); (3) the sparse reduced rank regression (SRRR) method as implemented in the R package rrpac [12], with pre-specified rank q and adaptive group LASSO penalty; (4) the sparse partial least squares (SPLS) approach as implemented in the R package spls [16], with the thresholding parameter η selected by CV, and the number of hidden components is set as q ; (5) the multivariate sparse group LASSO (MSGGLASSO) method as implemented via R package MSGLasso [30] with each predictor representing its own group; (6) the MLASSO method as implemented via the glmnet package [19] that penalizes the norm of each of the columns of the coefficient matrix; and (7) Multivariate square-root grouped LASSO (MSRL) with

each predictor representing its own group [34], with tuning parameter selected via 5-fold CV and the range of candidate tuning parameters, δ , set at 0.1.

For the frequentist procedures, the estimated model is defined to be the non-zero columns of the estimated coefficient matrix. For MBGL-SS the median thresholding estimator is used, and for MBSP the coefficients selected in the estimated model are those for which the 95% credible interval does not contain 0. For our EAS method we take the least square estimator of the MAP model as the point estimator for $\mathbf{B}^0_{M_o}$.

The metrics we use to evaluate the performance of the various methods, over 1,000 synthetic data sets for each of the six experiments in Table 1 and three experiments in Table 2, are the following. We report median MSPE on an out-of-sample test set, \mathbf{Y}_{new} , that is of the same size as \mathbf{Y} . We also report the average false discovery rate (FDR), the average false negative rate (FNR), average mis-classification probability (MP), average proportion of correct model selection (PCM), and median computation time. The results of the experiments in Table 1 are displayed in Table 3, and the results of the experiments in Table 2 are displayed in Table 4. The results are a bit less noisy than those reported in the simulation study in [1] because they only generated 100 synthetic datasets, and they did *not* report the out-of-sample prediction performance.

The explicit formulas for computing the metrics are, $\text{MSE} := \|\mathbf{Y} - \widehat{\mathbf{Y}}\|_{\mathbb{F}}^2/nq$, $\text{MSPE} := \|\mathbf{Y}_{new} - \widehat{\mathbf{Y}}_{new}\|_{\mathbb{F}}^2/nq$, $\text{FDR} := \text{FP}/(\text{FP}+\text{TP})$, $\text{FNR} := \text{FN}/(\text{FN}+\text{TN})$, and $\text{MP} := (\text{FP} + \text{FN})/pq$, where TP, FP, TN, and FN are, respectively, the number of true positives, false positives, true negatives, and false negatives. Moreover, we also present the average estimated posterior probability of the true model, denoted $\mathbb{P}(M_o | \mathbf{Y})$, for the Bayesian procedure and average fiducial probability, $r_\epsilon(M_o | \mathbf{Y})$, for the EAS procedure. Note that neither $r_\epsilon(M_o | \mathbf{Y})$ nor $\mathbb{P}(M_o | \mathbf{Y})$ can be calculated for the frequentist methods, or the MAP or credible region based Bayesian methods, like MBSP.

From Table 3, in the generic $n > p$ case for both sparsity levels of the true model, irrespective of whether ϵ is chosen based on BIC or CV, our EAS method performs on par with all other methods in terms of predictive performance, except SPLS which tends to exhibit inferior level of accuracy. In terms of variable selection performance, our EAS method chooses the correct model with a high probability and very low FDR and FNR, similar to the Bayesian methods MBGL-SS, MBSP and frequentist method SRRR. Other frequentist procedures tend to exhibit a lot of false positives. our EAS method does an excellent job in assigning a very high (GF) probability to the true model. Moreover, it is an advantage of our EAS method, and MBGL-SS, that they provide a probabilistic assessment of the competing models M so that inference can be made on how much better, say the MAP estimated model is from the second best model and so on. For instance, if there are many models that are assigned similar probabilities, then the practitioner is warned not to over-interpret inference based on a single model. This situation would possibly happen if there is sufficient collinearity (as defined by ϵ) among the important predictors. In that case, it is not reasonable to think that there is a unique choice of correct model (for fixed sample size) or

TABLE 3
Performance of EAS compared other methods over 1,000 replications.

Method	MSPE	FDR	FNR	MP	$\mathbb{P}(M_o \mathbf{Y})$	PCM	time (in sec)
LD ($n > p$), sparse : $n = 60, p = 30, q = 3, M_o = 5$							
EAS-BIC	2.17	0.0096	0	0.0007	0.952	0.952	143
EAS-CV	2.17	0.0148	0	0.0014	0.951	0.951	714
MBGL-SS	2.17	0.0042	0	0.0003	0.903	0.976	92
MBSP	2.28	0.0194	0	0.0013	N/A	0.891	20
MLASSO	2.47	0.6683	0	0.1245	N/A	0	0.6
MSGLASSO	2.65	0.4393	0	0.0553	N/A	0.032	0.6
SPLS	6.66	0.1419	0.023	0.0178	N/A	0.175	2.9
SRRR	2.17	0.0113	0	0.0008	N/A	0.939	0.1
MSRL	2.49	0.6603	0	0.1197	N/A	0	21
LD ($n > p$), dense : $n = 80, p = 60, q = 6, M_o = 40$							
EAS-BIC	4.14	0.0129	0	0.0021	0.9309	0.932	343
EAS-CV	4.12	0	0.0004	0	0.997	0.997	1944
MBGL-SS	4.10	0.0007	0	0.0001	0.9383	0.972	358
MBSP	4.25	0.0028	0	0.0003	N/A	0.897	30
MLASSO	5.94	0.331	0	0.055	N/A	0	1.1
MSGLASSO	6.23	0.3118	0	0.0505	N/A	0	3.0
SPLS	145.08	0.2738	0.2291	0.0467	N/A	0	7
SRRR	4.14	0.0033	0	0.0004	N/A	0.886	0.4
MSRL	5.69	0.3282	0	0.0543	N/A	0	32
HD ($p > n$), sparse : $n = 50, p = 200, q = 5, M_o = 20$							
EAS-BIC	3.40	0.0351	0.0002	0.0015	0.8798	0.882	368
EAS-CV	3.36	0.0075	0.0005	0.0004	0.9509	0.951	1811
MBGL-SS	52.29	0.5503	0.0076	0.0561	0.2029	0.232	3158
MBSP	4.43	0.0124	0	0.0003	N/A	0.778	95
MLASSO	15.85	0.7818	0.0001	0.0721	N/A	0	1.1
MSGLASSO	21.10	0.737	0.0049	0.0568	N/A	0	7.3
SPLS	112.05	0.5214	0.0401	0.0307	N/A	0	9.2
SRRR	15.54	0.7448	0.0016	0.0582	N/A	0	7.1
MSRL	16.69	0.7885	0.0002	0.0758	N/A	0	65
HD ($p > n$), dense : $n = 60, p = 100, q = 6, M_o = 40$							
EAS-BIC	6.91	0.0589	0.0135	0.0065	0.6673	0.669	441
EAS-CV	6.45	0.0203	0.0089	0.0026	0.8599	0.858	2252
MBGL-SS	6.31	0.0026	0.0002	0.0002	0.7697	0.936	861
MBSP	9.78	0.0297	0.0002	0.0021	N/A	0.339	56
MLASSO	34.86	0.5281	0.0006	0.0748	N/A	0	1.3
MSGLASSO	32.95	0.5265	0.001	0.0748	N/A	0	6.3
SPLS	197.30	0.4721	0.1437	0.0659	N/A	0	8.7
SRRR	22.03	0.474	0.0031	0.0603	N/A	0	2.7
MSRL	34.48	0.5511	0.001	0.0821	N/A	0	40
UHD ($p \gg n$), ultra-sparse : $n = 100, p = 500, q = 3, M_o = 10$							
EAS-BIC	2.23	0.0143	0	0.0004	0.9609	0.9609	241
EAS-CV	3.98	0.6135	0	0.0222	0.2116	0.2116	1645
MBGL-SS	2.23	0.0032	0	0.0001	0.7691	0.9729	18901
MBSP	2.86	0.0666	0	0.0005	N/A	0.5366	491
MLASSO	3.09	0.841	0	0.0396	N/A	0	1.3
MSGLASSO	16.71	0.7683	0.0013	0.0228	N/A	0	7.7
SPLS	27.54	0.2405	0.0055	0.0039	N/A	0.001	14
SRRR	5.91	0.9363	0	0.0981	N/A	0	20
MSRL	3.10	0.7890	0	0.0277	N/A	0	163
UHD ($p \gg n$), sparse : $n = 150, p = 1000, q = 4, M_o = 50$							
EAS-BIC	3.00	0	0	0	0.9944	0.995	1144
EAS-CV	3.00	0.0013	0	0.0001	0.9913	0.991	7095
MBGL-SS	354.19	0.932	0.0175	0.1532	0	0	116522
MBSP	3.46	0.0026	0	0	N/A	0.871	2316
MLASSO	19.00	0.8143	0	0.055	N/A	0	2.3
MSGLASSO	90.39	0.7995	0.0081	0.0455	N/A	0	74
SPLS	306.63	0.5796	0.0221	0.0193	N/A	0	91
SRRR	34.99	0.8052	0.0006	0.0514	N/A	0	88
MSRL	17.82	0.8323	0	0.0643	N/A	0	346

that the practitioner should choose a single model. For example, if two covariates are perfectly correlated, then the notion of *best subset* of the two covariates is not meaningful. Statistical inference guides data driven decisions, but such inference should also have the capability to suggest when there is not enough information in the data to make a decision; this embodies the EAS approach to model selection, and it also enables the EAS method to meaningfully be applied even in the absence of an underlying sparse data generating structure.

Continuing on Table 3, in the $p > n$ and $p \gg n$ scenarios, our EAS method does fulfill the expectations consistent with the strong model selection consistency. This is demonstrated by the fact that the EAS procedure, either using the BIC or CV tuning selection procedures, outperforms MBSP in terms of both the prediction performance and the average proportion of correct model selections. It also assigns a very high GF probability to the true model, which, again is consistent with our strong model selection consistency theoretical result.

The selection performance of the frequentist procedures seems to degrade rapidly when one moves from small to large p . They tend to select a lot of false signals in the estimated model. An interesting remark in support of MBGL-SS is that it seems to perform on par with the EAS procedure and outperform MBSP in some of the high-dimensional designs considered, arguably when either p is not so much larger than n or when the true model is ultra-sparse. This contradicts the numerical results presented in [1] where MBGL-SS is shown to perform poorly in all of the high-dimensional scenarios. In the $p \gg n$ and the ultra-sparse, when the optimal ϵ is chosen by CV, it seems to commit a lot of false discoveries. This might be attributed to the lack of identifiability of the true model in some of the folds, thus choosing an ϵ that it is relatively smaller than the optimal one. However, choosing ϵ through BIC seems to mitigate these computational bottlenecks observed with the CV procedure.

Table 4 compares performance of the methods with respect to challenges relating to larger dimension of the response, non-decaying correlation structure among the columns of the design matrix, and dense error covariance matrix. We do not present the results of MBGL-SS method because it takes more than 96 hours for a single dataset when q is increased to 60. Our results demonstrate that using either the BIC or CV tuning selection procedures, our EAS method outperforms all other methods in terms of both prediction performance and the average proportion of correct model selections. Furthermore, our EAS approach assigns a high GF probability to the true model, which aligns with our theoretical result regarding strong model selection consistency. In comparison to the frequentist approaches, SPLS is the only method that performs relatively well, whereas all other procedures tend to have a high FDR.

Lastly, observe that when the error covariance matrix is dense and the design matrix is complex—in addition to a large number of predictors—our EAS approach tends to result in higher false discoveries when optimal ϵ is selected by CV. As mentioned before, this issue may be due to a lack of identifiability of the true model in certain folds, resulting in the selection of a suboptimal ϵ that is relatively smaller than the optimal one. Nevertheless, selecting the regularization parameter using BIC appears to mitigate these computational bottlenecks

observed with the CV procedure. Overall, our findings demonstrate that our EAS method offers superior performance in these challenging scenarios, and may prove beneficial in various applied settings.

TABLE 4

Performance of EAS compared other methods for large dimension of the multivariate response, dense (non-decaying) correlation of the columns of the design matrix and error covariance matrix. Throughout this experiment, we consider UHD ($p \gg n$) and sparse case, i.e., the sample size, the number of predictors, and the size of the true model is fixed at $n = 150, p = 1000$, and $|M_o| = 50$.

Method	MSPE	FDR	FNR	MP	$\mathbb{P}(M_o \mathbf{Y})$	PCM	time (in sec)
Large dimension of response: $q = 60$ AR(1) correlation for the column of design matrix \mathbf{X} : $\Gamma_{ij} = 0.5^{ i-j }$ AR(1) error covariance matrix: $\mathbf{V}_{(M_o),ij}^0 = 2 \times 0.5^{ i-j }$							
EAS-BIC	3.01	0.000	0.000	0.000	0.99	1	4680
EAS-CV	3.01	0.000	0.000	0.000	0.99	1	22896
MBSP	3.79	0.054	0.000	0.000	N/A	0.09	4587
MLASSO	6.52	0.914	0.000	0.009	N/A	0.00	40
MSGLASSO	8.37	0.871	0.001	0.006	N/A	0.00	618
SPLS	3.43	0.167	0.000	0.0002	N/A	0	4280
SRRR	32.76	0.927	0.000	0.011	N/A	0	1922
MSRL	6.91	0.9334	0	0.0117	N/A	0	262
Large dimension of response: $q = 60$ Non-decaying correlation for the column of design matrix \mathbf{X} : $\Gamma_{ij} = 0.5\{1 + \mathbb{I}(i = j)\}$ AR(1) error covariance matrix: $\mathbf{V}_{(M_o),ij}^0 = 2 \times 0.5^{ i-j }$							
EAS-BIC	3.01	0.000	0.000	0.000	0.99	1	3174
EAS-CV	3.01	0.000	0.000	0.000	0.99	1	17680
MBSP	3.97	0.053	0.000	0.000	N/A	0.08	4521
MLASSO	7.98	0.496	0.000	0.001	N/A	0.00	36
MSGLASSO	6.56	0.78	0.000	0.003	N/A	0.00	7298
SPLS	5.14	0.185	0.000	0.0002	N/A	0	4799
SRRR	15.19	0.931	0.000	0.011	N/A	0	2848
MSRL	6.56	0.9406	0	0.0132	N/A	0	7098
Large dimension of response: $q = 60$ Non-decaying correlation for the column of design matrix \mathbf{X} : $\Gamma_{ij} = 0.5\{1 + \mathbb{I}(i = j)\}$ Dense error covariance matrix: $\mathbf{V}_{(M_o),ij}^0 = 1 + \mathbb{I}(i = j)$							
EAS-BIC	3.00	0.000	0.000	0.000	0.99	1	3214
EAS-CV	3.11	0.119	0.000	0.0002	0.72	0.72	17808
MBSP	2.98	0.052	0.000	0.000	N/A	0.12	4720
MLASSO	7.27	0.531	0.000	0.001	N/A	0.00	41
MSGLASSO	5.47	0.734	0.000	0.002	N/A	0.00	7184
SPLS	4.57	0.186	0.0001	0.0002	N/A	0	4800
SRRR	10.34	0.928	0.000	0.011	N/A	0	2876
MSRL	4.51	0.9374	0	0.0125	N/A	0	47963

6. Yeast cell data analysis

This section presents the results of implementing the EAS algorithm on the yeast cell cycle dataset [27]. The first part of the data contains yeast cell cycle gene expression data consisting of 18 measurements of messenger ribonucleic acid (mRNA) levels which are taken every 7 minutes of 119 minutes covering two cell cycle periods, for 542 cell cycle-related genes. The second part of the data contains binding information for a total of 106 transcription factors (TFs). TFs

are essential regulators of gene expression, and the binding of these proteins to specific DNA sequences controls the transcription of genes into mRNA, and ultimately the synthesis of functional proteins. The levels of mRNA, and therefore protein, produced by a gene are determined by the activity of the associated TFs [38, 51]. Therefore, it is important to identify the key TFs that regulate cell cycles. The particular dataset we analyze here is taken from the R package *spls* [16]. The response matrix \mathbf{Y} of gene-expression data is 18×542 dimensional and the design matrix \mathbf{X} of TFs is 106×542 dimensional, i.e., $n = 542$, $p = 106$, and $q = 18$. The data has been previously analyzed for variable selection of TFs in [1, 15].

As not all TFs are significantly contributing to the gene expressions; we aim to find a parsimonious set of cell-cycle regulating TFs by applying our EAS algorithm to this data. We use only the CV method to select the optimal tuning parameter ϵ for the real data application because we cannot rely on BIC due to the probable violation of the Gaussian assumption. We use a uniform grid of 16 values ranging from 0.01 to 0.2 for selection of ϵ via CV. We do not extend the grid endpoint over 0.2 because an implementation of Algorithm 3.2 did not select any TFs as admissible for epsilon greater than 0.2. We run the EAS method on the training set with the MCMC algorithm for 500 steps, discarding the initial 200 steps, and evaluate its performance on the validation set. The multivariate LASSO estimates serve as weights for proposing or removing predictors in the MCMC algorithm. We use the MAP estimated model from the MCMC sample as a point estimator to calculate the MSPE on the validation set. We select the optimal ϵ that minimizes the average MSPE across the 10 folds. The CV method selects optimal ϵ as 0.09. Finally, using $\epsilon = 0.09$, we implement Algorithm 3.2 on the entire yeast cell cycle data ten times to account for random variation in the MCMC chains, and each chain is run for 10,000 MCMC steps, discarding the initial 5,000 steps. Thus, we arrive at ten MCMC chains, each containing 5,000 samples from the GF distribution of the ϵ -admissible models for the yeast data.

The first five rows of Table 5 present the GF probabilities of the model, based on the ten MCMC chains. We truncate the table to display models with estimated GF probabilities greater than .1. The MAP estimated model identifies a rather parsimonious model, containing only 9 significant TFs. Interestingly, except for HIR1 and GAT3, all of the selected TFs in the MAP model are among the 21 experimentally confirmed cell-cycle-related TFs [52], supporting the relevance of our EAS approach to select a parsimonious model in real applications. Moreover, our research has identified HIR1 and GAT3 as two novel TFs. To evaluate the predictive performance of each of the models in Table 5, we employ ten-fold CV. Specifically, we use 90% of the data as the training set and obtain the least squares estimator. We then calculate the MSE and median absolute deviation (MAD) of the residuals on the remaining 10% of the data that we held out. We repeat this process 1,000 times, each time using different training and test sets, and compute the average MSE as MSPE and average MAD as mean absolute prediction error (MAPE). Finally, we scale the MSPE and MAPE by a factor of 100 for better clarity. The number of TFs in the optimally selected model for most of the competing methods are presented in Table 3 of [1]. We present the selected TFs in the optimal model for the Bayesian methods only,

TABLE 5
 Selected TFs for models with more than .1 estimated GF probability based on the EAS, MBSP and MBGL-SS methods, respectively. The optimal tuning parameter is selected via CV.

Method	Model	Selected TFs	# TFs	GF probability $r_e(M \mathbf{Y})$	MSPE	MAPE
EAS-CV	1	ACE2 GAT3 HIR1 MBP1 MCM1 NDD1 STE12 SWI5 SWI6	9	48.9%	19.0	23.9
	2	ACE2 GAT3 HIR1 MBP1 MCM1 NDD1 STE12 SWI5 SWI6 YAP5	10	18.9%	19.0	23.9
	3	ACE2 FKH2 GAT3 HIR1 MBP1 MCM1 NDD1 STE12 SWI4 SWI5	10	10%	19.0	23.6
MBSP		ACE2 FKH2 GAT3 HIR1 HIR2 MBP1 MET4 NDD1 REM1 STE12 SWI5 SWI6	12	NA	18.6	23.5
MBGL-SS		GAT3 NDD1 SWI5 SWI6	4	NA	20.1	23.6

i.e., MBSP and MBGL-SS, in the last two rows of Table 5. The last two columns of the table suggest that our EAS method does an outstanding job in terms of prediction accuracy in comparison to the competing methods while simultaneously maintaining parsimony in the selected model.

While models 2 and 3, each having 10 significant TFs, and the model selected by MBSP (with 12 significant TFs) exhibit similar or slightly better MSPE and MAPE compared to model 1, our EAS approach demonstrates remarkable performance by assigning a higher GF probability to the more parsimonious model 1. As determined by the EAS approach, the marginal improvement in prediction accuracy obtained with additional TFs is not substantial enough to justify the increase in model complexity.

The GF construction used enables parametric inference similar to Bayesian inference. In terms of quantifying uncertainty of the unknown parameters, the standard deviation of the estimated GF distribution of regression coefficient for each TFs, conditioned on a given model, plays a similar role to that of a standard error if the posterior mean is used as the point estimate. In this setting, we do not rely on p-values to determine statistical significance, but instead utilize the marginal inclusion probability of a specific TFs across all models in the GF distribution to establish its significance. The marginal inclusion probabilities of all TFs are presented in Table 6. We have shortened the table to only include TFs that have inclusion probabilities of at least .5. Remarkably, all the TFs with marginal inclusion probabilities exceeding .8 are the components of the MAP model listed in Table 5.

In this exposition, the real data analysis of yeast cell cycle data is not intended as a comprehensive investigation. Rather, it serves as a proof of concept for the practical utility of the EAS methodology for analyzing real data. It should be noted that obtaining a probability distribution of all possible models that are ϵ -admissible, as presented in Table 5, along with the marginal inclusion probabilities of the TFs, as in Table 6, is not feasible using frequentist or Bayesian point estimation-based procedures, such as MBSP. The ability of our

TABLE 6
Marginal inclusion probabilities of the TFs based on ten independent MCMC chains on the yeast cell cycle data. The optimal tuning parameter is selected via CV.

Transcription factors	Marginal inclusion probability (in %)
ACE2	100
HIR1	100
NDD1	100
STE12	100
SWI15	100
GAT3	99.9
MBP1	98.9
SWI16	95.0
MCM1	82.2
FKH2	58.5
RME1	53.5
HIR2	52.6
ARG81	50.0

EAS method to provide a probabilistic assessment of competing models stands out as an attractive feature of our EAS method. Although MCMC-based approaches are computationally more expensive, they provide more comprehensive information for uncertainty quantification.

7. Concluding remarks

The theoretical results presented in this article assume that the dimension q of response is fixed. However, we kept careful account of all instances of q in all of the non-asymptotic results presented, leaving an indication for the reader to understand the influence of q in the consistency rates. An obvious extension of our work is to allow q to grow, and in that case a careful account of the role of q in the theory is critical to determine the circumstances in which the EAS method remains a consistent model selection procedure. A field where this extended theory could be applied is functional data analysis (FDA), where the response is measured very densely for each subject and naturally the dimension of the response grows. Although smoothness in the mean and the covariance function is fundamental to the analysis of FDA, many FDA procedures simply rely on techniques that are developed for multivariate response data. Thus, we view this article as a promising first step on the pathway to a novel functional variable selection procedure.

Appendix A: Generalized fiducial distribution for multivariate linear regression

Here, we will derive the joint GF distribution of $(\mathbf{B}_{\cdot M} \mathbf{A}_{(M)})$ presented in (9). The data generating equation corresponding to i -th data, \mathbf{Y}_i is,

$$G_i := G(\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)}, \mathbf{U}_i) = \mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot, i} + \mathbf{A}_{(M)} \mathbf{U}_i \quad i = 1, 2, \dots, n,$$

where, $\mathbf{X}_{M \cdot, i}$ is the i th column of the matrix $\mathbf{X}_{M \cdot}$. Then,

$$\begin{aligned} \frac{\partial G_i}{\partial \mathbf{B}_{\cdot M}} &= \left[\frac{\partial(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot, i})}{\partial b_{11}}, \dots, \frac{\partial(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot, i})}{\partial b_{q|M|}} \right] \\ &= [\mathbf{J}^{11} \mathbf{X}_{M \cdot, i} \quad \dots \quad \mathbf{J}^{1|M|} \mathbf{X}_{M \cdot, i} \quad \dots \quad \mathbf{J}^{q1} \mathbf{X}_{M \cdot, i} \quad \dots \quad \mathbf{J}^{q|M|} \mathbf{X}_{M \cdot, i}] \\ &= \mathbf{I}_q \otimes \mathbf{X}_{M \cdot, i}^\top \in \mathbb{R}^{q \times q|M|}, \end{aligned}$$

where $\mathbf{J}^{kl} \in \mathbb{R}^{q \times |M|}$ is a sparse matrix of which (k, l) th element is 1 and all others are zero. Similarly, $\frac{\partial G_i}{\partial \mathbf{A}_{(M)}} = \mathbf{I}_q \otimes \mathbf{U}_i^\top \in \mathbb{R}^{q \times q^2}$. So the matrix of derivatives corresponding to the i -th data generating equation is,

$$\mathbf{D}_{(M), i} = \begin{bmatrix} \frac{\partial G_i}{\partial \mathbf{B}_{\cdot M}} & \vdots & \frac{\partial G_i}{\partial \mathbf{A}_{(M)}} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_{M \cdot, i}^\top & \vdots & \mathbf{I}_q \otimes \mathbf{U}_i^\top \end{bmatrix} \in \mathbb{R}^{q \times q|M|+q^2}.$$

The entire matrix of derivative corresponding to all observations is,

$$\mathbf{D}_{(M)} = \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_{M \cdot, 1}^\top & \mathbf{I}_q \otimes \mathbf{U}_1^\top \\ \vdots & \vdots \\ \mathbf{I}_q \otimes \mathbf{X}_{M \cdot, n}^\top & \mathbf{I}_q \otimes \mathbf{U}_n^\top \end{bmatrix} \in \mathbb{R}^{qn \times q|M|+q^2},$$

which is almost surely of full column rank if $n > |M| + q$. Now, define, $\mathbf{P}_{(M)} := [\mathbf{I}_q \otimes \mathbf{X}_{M \cdot}^\top \quad \mathbf{I}_q \otimes \mathbf{U}^\top]$, and evaluate the derivatives at $\mathbf{U} = \mathbf{A}_{(M)}^{-1} \tilde{\mathbf{U}}$ with $\tilde{\mathbf{U}} = \mathbf{Y} - \mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} \in \mathbb{R}^{q \times n}$. Because $\mathbf{P}_{(M)}$ is obtained rearranging rows of $\mathbf{D}_{(M)}$, $\mathbf{D}_{(M)}^\top \mathbf{D}_{(M)} = \mathbf{P}_{(M)}^\top \mathbf{P}_{(M)}$ and $\mathbf{P}_{(M)}^\top \mathbf{P}_{(M)}$ equal to

$$\begin{aligned} & \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_{M \cdot} \mathbf{X}_{M \cdot}^\top & \mathbf{I}_q \otimes \mathbf{X}_{M \cdot} \mathbf{U}^\top \\ \mathbf{I}_q \otimes \mathbf{U} \mathbf{X}_{M \cdot}^\top & \mathbf{I}_q \otimes \mathbf{U} \mathbf{U}^\top \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_{M \cdot} \mathbf{X}_{M \cdot}^\top & \mathbf{I}_q \otimes (\mathbf{X}_{M \cdot} \tilde{\mathbf{U}}^\top \mathbf{A}_{(M)}^{-1 \top}) \\ \mathbf{I}_q \otimes (\mathbf{A}_{(M)}^{-1} \tilde{\mathbf{U}} \mathbf{X}_{M \cdot}^\top) & \mathbf{I}_q \otimes \mathbf{A}_{(M)}^{-1} \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \mathbf{A}_{(M)}^{-1 \top} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{I}_{|M|} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_q \otimes \mathbf{A}_{(M)}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_{M \cdot} \mathbf{X}_{M \cdot}^\top & \mathbf{I}_q \otimes \mathbf{X}_{M \cdot} \tilde{\mathbf{U}}^\top \\ \mathbf{I}_q \otimes \tilde{\mathbf{U}} \mathbf{X}_{M \cdot}^\top & \mathbf{I}_q \otimes \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix} \\ & \times \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{I}_{|M|} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_q \otimes \mathbf{A}_{(M)}^{-1 \top} \end{bmatrix}. \end{aligned}$$

Then,

$$\det \mathbf{D}_{(M)}^\top \mathbf{D}_{(M)} = \left(\det \mathbf{A}_{(M)} \mathbf{A}_{(M)}^\top \right)^{-q} \det \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_M \cdot \mathbf{X}_M^\top & \mathbf{I}_q \otimes \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ \mathbf{I}_q \otimes \tilde{\mathbf{U}} \mathbf{X}_M^\top & \mathbf{I}_q \otimes \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix}.$$

After row and column operations,

$$\begin{aligned} \det \begin{bmatrix} \mathbf{I}_q \otimes \mathbf{X}_M \cdot \mathbf{X}_M^\top & \mathbf{I}_q \otimes \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ \mathbf{I}_q \otimes \tilde{\mathbf{U}} \mathbf{X}_M^\top & \mathbf{I}_q \otimes \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix} &= \det \left\{ \mathbf{I}_q \otimes \begin{bmatrix} \mathbf{X}_M \cdot \mathbf{X}_M^\top & \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ \tilde{\mathbf{U}} \mathbf{X}_M^\top & \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix} \right\} \\ &= \left(\det \begin{bmatrix} \mathbf{X}_M \cdot \mathbf{X}_M^\top & \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ \tilde{\mathbf{U}} \mathbf{X}_M^\top & \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix} \right)^q. \end{aligned}$$

Further, by property of determinant of block matrices,

$$\begin{aligned} \det \begin{bmatrix} \mathbf{X}_M \cdot \mathbf{X}_M^\top & \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ \tilde{\mathbf{U}} \mathbf{X}_M^\top & \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \end{bmatrix} &= (\det \mathbf{X}_M \cdot \mathbf{X}_M^\top) \times \\ &\det \left[\tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top - \tilde{\mathbf{U}} \mathbf{X}_M^\top (\mathbf{X}_M \cdot \mathbf{X}_M^\top)^{-1} \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \right]. \end{aligned}$$

Analogous to univariate linear regression,

$$\mathbf{Y} - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M = \mathbf{Y} - \mathbf{Y} \mathbf{X}_M^\top (\mathbf{X}_M \cdot \mathbf{X}_M^\top)^{-1} \mathbf{X}_M = \mathbf{Y} (\mathbf{I} - \mathbf{H}_{(M)}),$$

and,

$$\begin{aligned} &\tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top - \tilde{\mathbf{U}} \mathbf{X}_M^\top (\mathbf{X}_M \cdot \mathbf{X}_M^\top)^{-1} \mathbf{X}_M \cdot \tilde{\mathbf{U}}^\top \\ &= \left(\tilde{\mathbf{U}} - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M + \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M \right) [\mathbf{I} - \mathbf{H}_{(M)}] \left(\tilde{\mathbf{U}} - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M + \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M \right)^\top \\ &= \left(\mathbf{Y} - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M \right) [\mathbf{I} - \mathbf{H}_{(M)}] \left(\mathbf{Y} - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M \right)^\top \\ &= \hat{\Sigma}_{(M)}. \end{aligned}$$

The Jacobian is obtained as

$$J(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) := \sqrt{\det \mathbf{D}_{(M)}^\top \mathbf{D}_{(M)}} = C_M (\det \mathbf{V}_{(M)})^{-q/2},$$

where $C_M := (\det \mathbf{\Omega}_{(M)})^{q/2} (\det \hat{\Sigma}_{(M)})^{q/2}$ and $\mathbf{\Omega}_{(M)} = \mathbf{X}_M \cdot \mathbf{X}_M^\top$. Note that C_M does not depend on either of $\mathbf{B}_{\cdot M}$ or $\mathbf{A}_{(M)}$. The likelihood function is,

$$f(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) \propto (\det \mathbf{V}_{(M)})^{-\frac{n}{2}} \exp \left[-\frac{1}{2} \text{tr} \left(\mathbf{R}_{(M)} \mathbf{V}_{(M)}^{-1} \right) \right],$$

where,

$$\begin{aligned}\mathbf{R}_{(M)} &= (\mathbf{Y} - \mathbf{B}_{\cdot M} \mathbf{X}_{M\cdot}) (\mathbf{Y} - \mathbf{B}_{\cdot M} \mathbf{X}_{M\cdot})^\top \\ &= \widehat{\boldsymbol{\Sigma}}_{(M)} + \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right) \mathbf{X}_{M\cdot} \mathbf{X}_{M\cdot}^\top \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right)^\top.\end{aligned}$$

Under the matrix change of variables $\mathbf{V}_{(M)} = \mathbf{A}_{(M)} \mathbf{A}_{(M)}^\top$, we use Theorem 1.4.10 of [23], to derive the marginal GF distribution of model M as,

$$\begin{aligned}r_\epsilon(M|\mathbf{Y}) &\propto \int \int f(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) J(\mathbf{Y}, (\mathbf{B}_{\cdot M}, \mathbf{A}_{(M)})) h_\epsilon(\mathbf{B}_{\cdot M}) d(\mathbf{A}_{(M)}, \mathbf{B}_{\cdot M}) \\ &\propto C_M \int_{\mathbf{B}_{\cdot M}} \int_{\mathbf{A}_{(M)}} (\det \mathbf{V}_{(M)})^{-\frac{q+n}{2}} e^{-\frac{1}{2} \text{tr}(\mathbf{R}_{(M)} \mathbf{V}_{(M)}^{-1})} h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{A}_{(M)} d\mathbf{B}_{\cdot M} \\ &\propto C_M \int_{\mathbf{B}_{\cdot M}} \int_{\mathbf{V}_{(M)} > \mathbf{0}} (\det \mathbf{V}_{(M)})^{-\frac{q+n+1}{2}} e^{-\frac{1}{2} \text{tr}(\mathbf{R}_{(M)} \mathbf{V}_{(M)}^{-1})} h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{V}_{(M)} d\mathbf{B}_{\cdot M} \\ &\propto C_M \int_{\mathbf{B}_{\cdot M}} [\det \mathbf{R}_{(M)}]^{-\frac{n}{2}} h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{B}_{\cdot M} \quad (\text{w.r.t inverse-wishart kernel}) \\ &= C_M \int \left(\det \left[\widehat{\boldsymbol{\Sigma}}_{(M)} + \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right) \boldsymbol{\Omega}_{(M)} \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right)^\top \right] \right)^{-\frac{n}{2}} h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{B}_{\cdot M} \\ &= C_M \int \left(\det \left[\mathbf{I}_M + \boldsymbol{\Omega}_{(M)} \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right)^\top \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1} \left(\mathbf{B}_{\cdot M} - \widehat{\mathbf{B}}_{\cdot M} \right) \right] \det \widehat{\boldsymbol{\Sigma}}_{(M)} \right)^{-\frac{n}{2}} \\ &\quad \times h_\epsilon(\mathbf{B}_{\cdot M}) d\mathbf{B}_{\cdot M} \\ &\propto \Gamma_q \left(\frac{n - |M|}{2} \right) \pi^{\frac{q|M|}{2}} \left(\det \widehat{\boldsymbol{\Sigma}}_{(M)} \right)^{-\left(\frac{n - |M| - q}{2} \right)} \mathbb{E} (h_\epsilon(\mathbf{B}_{\cdot M})),\end{aligned}$$

where the last line is obtained by integrating with respect to matrix-t kernel. This completes the derivation of the GF distribution of $(\mathbf{B}_{\cdot M} \mathbf{A}_{(M)})$ for multivariate linear regression setting under the model (1).

Appendix B: Proof of lemmas and theorems

Proof of Lemma 4.2.

$$\begin{aligned}\lambda_v^{-1} \lambda_{\min}(\widehat{\boldsymbol{\Sigma}}_{(M)}) &= \|\mathbf{A}_{(M_o)}^{0^{-1}}\|^2 \lambda_{\min}(\mathbf{Y}(\mathbf{I}_n - \mathbf{H}_{(M)})\mathbf{Y}^\top) \\ &\geq \lambda_{\min}(\mathbf{A}_{(M_o)}^{0^{-1}} \mathbf{Y}(\mathbf{I}_n - \mathbf{H}_{(M)})\mathbf{Y}^\top \mathbf{A}_{(M_o)}^{0^{-\top}}) \\ &= \lambda_{\min}(\mathbf{Z}(\mathbf{I}_n - \mathbf{H}_{(M)})\mathbf{Z}^\top)\end{aligned}$$

where $\mathbf{Z} := \mathbf{A}_{(M_o)}^{0^{-1}} \mathbf{Y} = \mathbb{E}(\mathbf{Z}) + \mathbf{U}$, with $\mathbb{E}(\mathbf{Z}) = \mathbf{A}_{(M_o)}^{0^{-1}} \mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o\cdot}$. Recall that $\mathbf{H}_{(M)}$ is the symmetric projection matrix onto row space of $\mathbf{X}_{M\cdot}$, and let $r := \text{rank}(\mathbf{H}_{(M)})$. Since $\mathbf{I}_n - \mathbf{H}_{(M)}$ is symmetric and idempotent with rank $n - r$, there exists a $\mathbf{G}_{(M)} \in \mathbb{R}^{n \times (n-r)}$ such that $\mathbf{G}_{(M)} \mathbf{G}_{(M)}^\top = (\mathbf{I} - \mathbf{H}_{(M)})$ and $\mathbf{G}_{(M)}^\top \mathbf{G}_{(M)} = \mathbf{I}_{n-r}$. Define $\widetilde{\mathbf{Z}}_M := \mathbf{Z} \mathbf{G}_{(M)} = \mathbb{E}(\mathbf{Z} \mathbf{G}_{(M)}) + \mathbf{U} \mathbf{G}_{(M)}$. Then, $\mathbf{Z}(\mathbf{I}_n - \mathbf{H}_{(M)})\mathbf{Z}^\top = \widetilde{\mathbf{Z}}_{(M)} \widetilde{\mathbf{Z}}_{(M)}^\top$ and the minimum eigenvalue of $\widetilde{\mathbf{Z}}_{(M)} \widetilde{\mathbf{Z}}_{(M)}^\top$ will

be the minimum singular value of $\tilde{\mathbf{Z}}_{(M)}$ (in the compact SVD notation). Notice that, $(\mathbf{U}\mathbf{G}_{(M)})^\top \sim \text{Matrix-Normal}_{n-r,q}(\mathbf{0}, \mathbf{I}_{n-r}, \mathbf{I}_q)$. Applying Theorem 2.1 of [44] for independent non-centered Gaussian design for sufficiently large n with $t = \sqrt{n-r-q} - \sqrt{q} - 1 - \sqrt{\tau}\sqrt{n-r} > 0$, $n-r > 2q$, and $0 < \tau < 1$, we get

$$\mathbb{P}_y \left(\sigma_{\min}(\tilde{\mathbf{Z}}_{(M)}) \geq \sqrt{\tau(n-r)} \right) > 1 - e^{-\{\sqrt{n-r-q} - \sqrt{q} - 1 - \sqrt{\tau}\sqrt{n-r}\}^2}.$$

Since, $(\sqrt{n-r-q} - \sqrt{q} - 1 - \sqrt{\tau}\sqrt{n-r}) / (1 - \sqrt{\tau})\sqrt{n-r} \rightarrow 1$ as $n \rightarrow \infty$, for sufficiently large n , the exponent in the right hand of the above expression will be larger than $(1 - \sqrt{\tau})\sqrt{n-r}/\sqrt{2}$. The proof is completed by observing that $n-r \geq n - |M|$ and $\lambda_v > 0$, by Condition 4.1. \square

Proof of Theorem 4.5. Case 1. First we consider models M with $M \subsetneq M_o$. Let $r_{n,M} = 1 - (1 + \xi_{n,|M|}\lambda_v/\bar{\lambda}_v)/2$. Then for sufficiently large n , $r_{n,M} \in (0, 1/2)$. We take $\tau = r_{n,M}$ in Lemma 4.2 to get,

$$\mathbb{P}_y \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) < r_{n,M}\lambda_v(n - |M|) \right) \leq e^{-(1-1/\sqrt{2})^2(n-|M|)/2} < e^{-0.04(n-|M|)},$$

Then,

$$\begin{aligned} & \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \xi_{n,|M|}^q \right) \\ & \leq \mathbb{P}_y \left(\det \widehat{\Sigma}_{(M_o)} \geq \xi_{n,|M|}^q \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) \right)^q \right) \\ & \leq \mathbb{P}_y \left(\det \widehat{\Sigma}_{(M_o)} \geq \xi_{n,|M|}^q (r_{n,M}\lambda_v(n - |M|))^q \right) \\ & \quad + \mathbb{P} \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) < r_{n,M}\lambda_v(n - |M|) \right) \\ & \leq \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \mathbf{V}_{(M_o)}^0} \geq \xi_{n,|M|}^q \left(\frac{\lambda_v}{\bar{\lambda}_v} \right)^q (r_{n,M}(n - |M|))^q \right) + e^{-0.04(n-|M|)}, \end{aligned}$$

and by Theorem 3.3.22 of [23] with $\xi_{n,|M|}^* = \xi_{n,|M|}\lambda_v/\bar{\lambda}_v \in (0, 1)$,

$$\begin{aligned} & \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \xi_{n,|M|}^q \right) \\ & \leq \sum_{i=1}^q \mathbb{P}_y \left(\chi_{n-|M_o|-i+1}^2 \geq r_{n,M}\xi_{n,|M|}^*(n - |M|) \right) + e^{-0.04(n-|M|)} \\ & \leq \sum_{i=1}^q e^{-s r_{n,M}\xi_{n,|M|}^*(n-|M|) - \left(\frac{n-|M_o|-i+1}{2}\right) \log(1-2s)} + e^{-0.04(n-|M|)} \\ & \leq q \exp \left(-(n - |M|) \left(s r_{n,M}\xi_{n,|M|}^* + \frac{1}{4} \log(1 - 2s) \right) \right) + e^{-0.04(n-|M|)}, \end{aligned}$$

where the second to last inequality holds for any $s < 1/2$ by Chernoff's bound, and the last inequality holds for any $n > 2|M_o| + 2q - |M| \geq 2q + 1$. Taking

$$\begin{aligned}
 s &= \frac{1}{2} \left(1 - \frac{1}{2r_{n,M} \xi_{n,|M|}^*} \right) < 0, \\
 & s r_{n,M} \xi_{n,|M|}^* + \frac{1}{4} \log(1 - 2s) \\
 &= \frac{1}{4} \left(\xi_{n,|M|}^* (1 - \xi_{n,|M|}^*) - 1 - \log \left(\xi_{n,|M|}^* (1 - \xi_{n,|M|}^*) \right) \right) \\
 &\geq \frac{1}{4} \left(\xi_{n,|M|}^* (1 - \xi_{n,|M|}^*) - 1 + \log 4 \right) \geq \frac{\xi_{n,|M|} \lambda_v n^\alpha \log |M_o|}{2\lambda_v (n - |M|)} + 0.09 > 0,
 \end{aligned}$$

which implies that,

$$\mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \xi_{n,|M|}^q \right) \leq qe^{-\frac{\xi_{n,|M|} n^\alpha \log |M_o| \lambda_v}{2\lambda_v} - 0.09(n - |M|)} + e^{-0.04(n - |M|)}.$$

Therefore,

$$\begin{aligned}
 & \mathbb{P}_y \left(\bigcup_{M \subsetneq M_o} \left\{ \mathbf{Y} : \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right)^{\frac{n - |M| - q}{2}} > \xi_{n,|M|}^{\frac{q(n - |M| - q)}{2}} \right\} \right) \\
 &\leq \sum_{j=1}^{M_o} \binom{|M_o|}{j} \max_{M \subsetneq M_o: |M|=j} \mathbb{P}_y \left(\mathbf{Y} : \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right)^{\frac{n - j - q}{2}} > \xi_{n,j}^{\frac{q(n - j - q)}{2}} \right) \\
 &\leq \sum_{j=1}^{M_o} e^{j \log(|M_o|)} \max_{M \subsetneq M_o: |M|=j} \left\{ qe^{-\frac{\xi_{n,j} n^\alpha \log |M_o| \lambda_v}{2\lambda_v} - 0.09(n - j)} + e^{-0.04(n - j)} \right\} \\
 &\leq |M_o| q \exp \left(-\frac{\xi_{n,|M_o|} n^\alpha \log |M_o| \lambda_v}{2\lambda_v} - 0.09(n - |M_o|) + |M_o| \log(|M_o|) \right) \\
 &\quad + 2 |M_o| \exp(-0.04(n - |M_o|) + |M_o| \log |M_o|),
 \end{aligned}$$

The proof for Case 1 completes by noting that,

$$\begin{aligned}
 \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right)^{\frac{n - |M| - q}{2}} &\leq \xi_{n,|M|}^{\frac{q(n - |M| - q)}{2}} \\
 &\leq \left(1 - \frac{2n^\alpha \log |M_o|}{n - |M| - q} \right)^{\frac{q(n - |M| - q)}{2}} \\
 &\leq e^{-qn^\alpha \log |M_o|}.
 \end{aligned}$$

Case 2. Fix an arbitrary model M such that $M \not\subseteq M_o$ and $|M| \leq n^\alpha$, and construct a new model $M' := M \cup M_o$ where $|M'| = |M| + \ell$ for some $\ell \in \{1, \dots, |M_o|\}$. Further, $M' \supset M_o$ implies $(\mathbf{H}_{(M')} - \mathbf{H}_{(M_o)}) (\mathbf{I} - \mathbf{H}_{(M')}) = \mathbf{0}$, which means $\widehat{\Sigma}_{(M_o)} - \widehat{\Sigma}_{(M')}$ is independent of $\widehat{\Sigma}_{(M')}$. For sufficiently large n ,

$$\widehat{\Sigma}_{(M')} \sim \text{Wishart}_q \left(n - r - \ell, \mathbf{V}_{(M_o)}^0 \right)$$

$$\widehat{\Sigma}_{(M_o)} - \widehat{\Sigma}_{(M')} \sim \text{Wishart}_q \left(r + \ell - |M_o|, \mathbf{V}_{(M_o)}^0 \right),$$

where $r := \text{rank}(\mathbf{X}_{M'}) \leq |M|$. Without loss of generality for the bounds we derive it suffices to work with $r = |M|$. By Theorem 10.5.3 of [35],

$$\frac{\det \widehat{\Sigma}_{(M')}}{\det \widehat{\Sigma}_{(M_o)}} \sim \begin{cases} \prod_{i=1}^q V_i^{(1)} & \text{if } |M| + \ell - |M_o| \geq q \\ \prod_{i=1}^{|M|+\ell-|M_o|} V_i^{(2)} & \text{if } |M| + \ell - |M_o| < q \end{cases}$$

where

$$V_i^{(1)} \stackrel{\text{ind}}{\sim} \text{Beta} \left(\frac{1}{2} (n - |M| - \ell - i + 1), \frac{1}{2} (|M| + \ell - |M_o|) \right),$$

$$V_i^{(2)} \stackrel{\text{ind}}{\sim} \text{Beta} \left(\frac{1}{2} (n - |M_o| - q - i + 1), \frac{q}{2} \right).$$

We will handle the two cases separately.

Case 2a: Suppose $|M| + \ell - |M_o| \geq q$. Because, $M' \supset M$, $\widehat{\Sigma}_{(M)} - \widehat{\Sigma}_{(M')}$ is a positive definite matrix and that implies $\det \widehat{\Sigma}_{(M)} \geq \det \widehat{\Sigma}_{(M')}$. Hence,

$$\begin{aligned} \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \zeta_{n,|M|}^q \right) &\leq \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M')}} > \zeta_{n,|M|}^q \right) \\ &\leq \sum_{i=1}^q \mathbb{P}_y \left(V_i^{(1)} < \zeta_{n,|M|}^{-1} \right). \end{aligned}$$

The probabilities in the sum on the right side are bounded by approximating the CDF of the Beta density, and approximating of the beta function by Theorem 2 of [25] as,

$$\begin{aligned} &\mathbb{P}_y \left(V_i^{(1)} < \zeta_{n,|M|}^{-1} \right) \\ &= \left(B \left(\frac{n - |M| - \ell - i + 1}{2}, \frac{|M| + \ell - |M_o|}{2} \right) \right)^{-1} \\ &\quad \times \int_0^{\zeta_{n,|M|}^{-1}} y^{\frac{1}{2}(n - |M| - \ell - i + 1) - 1} (1 - y)^{\frac{1}{2}(|M| + \ell - |M_o|) - 1} dy \\ &\leq \frac{\left(\frac{n - |M_o| - i + 1}{2} \right)^{\left(\frac{|M| + \ell - |M_o|}{2} \right)}}{\Gamma \left(\frac{|M| + \ell - |M_o|}{2} \right)} \zeta_{n,|M|}^{-\left(\frac{n - |M| - \ell - i + 1}{2} \right)} \left[1 - \left(\frac{\zeta_{n,|M|} - 1}{\zeta_{n,|M|}} \right)^{\frac{|M| + \ell - |M_o|}{2}} \right] \\ &\leq \left(\frac{n - |M_o|}{2} \right)^{\frac{|M|}{2}} \zeta_{n,|M|}^{-\left(\frac{n - |M| - q}{2} \right) + \frac{|M_o| + 1}{2}} \cdot 1 \end{aligned}$$

Case 2b: Suppose $|M| + \ell - |M_o| < q$. Then similar to the previous case,

$$\mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \zeta_{n,|M|}^q \right) \leq \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M')}} > \zeta_{n,|M|}^{|M| + \ell - |M_o|} \right)$$

$$\leq \sum_{i=1}^{|M|+\ell-|M_o|} \mathbb{P}_y \left(V_i^{(2)} < \zeta_{n,|M|}^{-1} \right),$$

and,

$$\begin{aligned} & \mathbb{P}_y \left(V_i^{(2)} < \zeta_{n,|M|}^{-1} \right) \\ &= \frac{1}{B\left(\frac{n-|M_o|-q-i+1}{2}, \frac{q}{2}\right)} \int_0^{\zeta_{n,|M|}^{-1}} y^{\frac{1}{2}(n-|M_o|-q-i+1)-1} (1-y)^{\frac{q}{2}-1} dy \\ &\leq \frac{\left(\frac{n-|M_o|-i+1}{2}\right)^{\frac{q}{2}}}{\Gamma\left(\frac{q}{2}\right)} \zeta_{n,|M|}^{-\left(\frac{n-|M_o|-q-i-1}{2}\right)} \left[1 - \left(\frac{\zeta_{n,|M|}-1}{\zeta_{n,|M|}}\right)^{\frac{q}{2}} \right] \\ &\leq \left(\frac{n-|M_o|}{2}\right)^{\frac{q}{2}} \zeta_{n,|M|}^{-\left(\frac{n-|M|-q}{2}\right)+\frac{|M_o|+1}{2}}. \end{aligned}$$

Thus, in any case, for sufficiently large n ,

$$\begin{aligned} & \mathbb{P}_y \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} > \zeta_{n,|M|}^q \right) \\ &\leq 2q \exp \left(-n^\alpha \log(n - |M_o|) + \max \left\{ \frac{|M|}{2}, \frac{q}{2} \right\} \log \left(\frac{n - |M_o|}{2} \right) \right. \\ &\quad \left. - |M| \log p + \frac{(|M_o| + 1)}{2} \log(\zeta_{n,|M|}) \right). \end{aligned}$$

Therefore,

$$\begin{aligned} & \mathbb{P}_y \left(\bigcup_{M \not\subseteq M_o: |M| \leq n^\alpha} \left\{ \mathbf{Y} : \left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right)^{\frac{n-|M|-q}{2}} > \zeta_{n,|M|}^{\frac{q(n-|M|-q)}{2}} \right\} \right) \\ &\leq \sum_{j=1}^{n^\alpha} \binom{p}{j} \max_{M \not\subseteq M_o: |M|=j} \mathbb{P}_y \left(\left(\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right)^{\frac{n-j-q}{2}} > \zeta_{n,j}^{\frac{q(n-j-q)}{2}} \right) \\ &\leq \sum_{j=1}^{n^\alpha} \left\{ 2q \exp \left(-n^\alpha \log(n - |M_o|) + \max \left\{ \frac{j}{2}, \frac{q}{2} \right\} \log(n - |M_o|) \right. \right. \\ &\quad \left. \left. + \frac{(|M_o| + 1)}{2} \log(\zeta_{n,j}) \right) \right\} \\ &\leq 2q \exp \left(-\frac{n^\alpha}{2} \log(n - |M_o|) + \alpha \log n + \frac{(|M_o| + 1)}{2} \log(\zeta_{n,n^\alpha}) \right). \end{aligned}$$

The proof is completed by noting that $\zeta_{n,|M|}^{\frac{q(n-|M|-q)}{2}} \leq e^{q(n^\alpha \log(n-|M_o|)+|M| \log p)}$. □

Proof of Theorem 4.7. To prove Theorem 4.7 and Theorem 4.9 we need two additional results, stated below. The proofs of Lemma B.1 and Lemma B.2 are provided after the proof of the Theorem 4.9. Since these two lemmas are only necessary to bound the $E(h_\epsilon(\mathbf{B}_{\cdot M}))$ from above for the large models, without loss of generality we assume that $\mathbf{X}_{M\cdot}$ is of full row rank, as in the statement of both the lemmas. If $\mathbf{X}_{M\cdot}$ is not of full rank, $h_\epsilon(\mathbf{B}_{\cdot M})$ is automatically zero.

Lemma B.1. *For model M with $|M| < n - 4q$,*

$$\mathbf{B}_{\cdot M} \sim T_{q,M} \left(n - |M| - q + 1, \widehat{\mathbf{B}}_{\cdot M}, \widehat{\boldsymbol{\Sigma}}_{(M)}, (\mathbf{X}_{M\cdot} \mathbf{X}_{M\cdot}^\top)^{-1} \right),$$

where $\widehat{\mathbf{B}}_{\cdot M} = \mathbf{Y} \mathbf{X}_{M\cdot}^\top (\mathbf{X}_{M\cdot} \mathbf{X}_{M\cdot}^\top)^{-1}$, and for any $\epsilon > 0$,

$$\mathbb{P} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M\cdot} - \widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_{M\cdot} \right) \right\|_F^2 \leq \frac{\epsilon}{9} \right) \geq 1 - V_{5,n,M},$$

with,

$$V_{5,n,M} := \exp \left(-\frac{\epsilon(n - |M|)}{36} + \frac{q|M|}{2} \right) + 2 \exp \left(-\frac{(\sqrt{n - |M|} - 2\sqrt{q})^2}{8} \right).$$

Lemma B.2. *For any model M such that $|M| < n - 4q$, then the least-squared estimator,*

$$\widehat{\mathbf{B}}_{\cdot M} \sim \text{Matrix-Normal}_{q,|M|} \left(\mathbb{E}_y \left(\widehat{\mathbf{B}}_{\cdot M} \right), \mathbf{V}_{(M_o)}^0, (\mathbf{X}_{M\cdot} \mathbf{X}_{M\cdot}^\top)^{-1} \right),$$

and for sufficiently large n ,

$$\mathbb{P}_y \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left[\widehat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\widehat{\mathbf{B}}_{\cdot M} \right) \right] \mathbf{X}_{M\cdot} \right\|_F^2 \leq \frac{\epsilon}{9} \right) \geq 1 - V_{6,n,M},$$

where,

$$V_{6,n,M} := \exp \left(-\frac{\epsilon(n - |M|) \underline{\lambda}_v}{36 \bar{\lambda}_v} + \frac{q|M|}{2} \right) + \exp(-0.04(n - |M|)).$$

Let $\widetilde{\mathbf{B}}_{\min}$ minimizes $\frac{1}{2} \left\| \mathbf{A}_{(M_o)}^{0^{-1}} (\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o\cdot} - \mathbf{B} \mathbf{X}) \right\|_F^2$ subject to the constraint $|\{j : \|\mathbf{B}_j\| \neq 0\}|$ less than equal to $|M_o| - 1$. Also, suppose \mathbf{B}_{\min} minimizes $\left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\mathbf{B}_{\cdot M_o} \mathbf{X}_{M_o\cdot} - \mathbf{B} \mathbf{X}) \right\|_F^2$ subject to $|\{j : \|\mathbf{B}_j\| \neq 0\}| \leq |M_o| - 1$. Then,

$$\begin{aligned} & \left\| \mathbf{A}_{(M_o)}^{0^{-1}} \left(\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o\cdot} - \widetilde{\mathbf{B}}_{\min} \mathbf{X} \right) \right\|_F \\ & \leq \left\| \mathbf{A}_{(M_o)}^{0^{-1}} \left(\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o\cdot} - \mathbf{B}_{\min} \mathbf{X} \right) \right\|_F \end{aligned}$$

$$\leq \left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbb{F}} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o \cdot} - \mathbf{B}_{\min} \mathbf{X}) \right\|$$

Since $\mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)} \mathbf{A}_{(M_o)}^{0^{-T}} \sim \text{Wishart}_q(n - |M_o|, \mathbf{I}_q)$, $\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbb{F}}^2$ follows a chi-squared distribution with $q(n - |M_o|)$ degrees of freedom. Using triangle inequality,

$$\begin{aligned} & \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o \cdot} - \mathbf{B}_{\min} \mathbf{X}) \right\|_{\mathbb{F}} \\ & \leq \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} [\mathbf{B}_{\cdot M_o} - \widehat{\mathbf{B}}_{\cdot M_o}] \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}} + \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\widehat{\mathbf{B}}_{\cdot M_o} - \mathbf{B}_{\cdot M_o}^0) \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}} \\ & \quad + \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\mathbf{B}_{\cdot M_o} \mathbf{X}_{M_o \cdot} - \mathbf{B}_{\min} \mathbf{X}) \right\|_{\mathbb{F}}. \end{aligned} \tag{12}$$

Then,

$$\begin{aligned} & \mathbf{I} \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} (\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o \cdot} - \widetilde{\mathbf{B}}_{\min} \mathbf{X}) \right\|_{\mathbb{F}}^2 > 36q(n - |M_o|)\epsilon \right) \\ & \leq \mathbf{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} [\mathbf{B}_{\cdot M_o} - \widehat{\mathbf{B}}_{\cdot M_o}] \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}}^2 > \epsilon \right) + \\ & \quad \mathbf{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\widehat{\mathbf{B}}_{\cdot M_o} - \mathbf{B}_{\cdot M_o}^0) \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}}^2 > \epsilon \right) + \\ & \quad \mathbf{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\mathbf{B}_{\cdot M_o} \mathbf{X}_{M_o \cdot} - \mathbf{B}_{\min} \mathbf{X}) \right\|_{\mathbb{F}}^2 > \epsilon \right) + \\ & \quad \mathbf{I} \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbb{F}}^2 > 2q(n - |M_o|) \right). \end{aligned}$$

Taking expectation with respect to Fiducial distribution of $\mathbf{B}_{\cdot M_o}$ given \mathbf{Y} ,

$$\begin{aligned} & \mathbf{I} \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} (\mathbf{B}_{\cdot M_o}^0 \mathbf{X}_{M_o \cdot} - \widetilde{\mathbf{B}}_{\min} \mathbf{X}) \right\|_{\mathbb{F}}^2 > 36q(n - |M_o|)\epsilon \right) \\ & \leq \mathbb{P} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} [\mathbf{B}_{\cdot M_o} - \widehat{\mathbf{B}}_{\cdot M_o}] \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}}^2 > \epsilon \right) \\ & \quad + \mathbf{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\widehat{\mathbf{B}}_{\cdot M_o} - \mathbf{B}_{\cdot M_o}^0) \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}}^2 > \epsilon \right) \end{aligned} \tag{13}$$

$$+ \mathbb{E}(h_\epsilon(\mathbf{B}_{\cdot M_o})) + \mathbf{I} \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbb{F}}^2 > 2q(n - |M_o|) \right). \tag{14}$$

Next,

$$\mathbb{P}_y \left(\left\{ \mathbf{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} (\widehat{\mathbf{B}}_{\cdot M_o} - \mathbf{B}_{\cdot M_o}^0) \mathbf{X}_{M_o \cdot} \right\|_{\mathbb{F}}^2 > \epsilon \right) = 0 \right\} \cap \right.$$

$$\begin{aligned}
 & \left\{ \mathbf{I} \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbf{F}}^2 > 2q(n - |M_o|) \right) = 0 \right\} \\
 & \geq \mathbb{P}_y \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{-1/2} \left(\widehat{\mathbf{B}}_{\cdot M_o} - \mathbf{B}_{\cdot M_o}^0 \right) \mathbf{X}_{M_o} \right\|_{\mathbf{F}}^2 \leq \epsilon \right) \\
 & \quad + \mathbb{P}_y \left(\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \widehat{\boldsymbol{\Sigma}}_{(M_o)}^{1/2} \right\|_{\mathbf{F}}^2 \leq 2q(n - |M_o|) \right) - 1 \\
 & \geq 1 - e^{-\frac{\epsilon(n - |M_o|)\lambda_v}{4\lambda_v} + \frac{q|M_o|}{2}} - \exp(-0.04(n - |M_o|)) - e^{-0.15q(n - |M_o|)},
 \end{aligned}$$

where the first probability is obtained by an application of Lemma B.2 and the second probability is computed by the Chernoff bound for the χ^2 distribution with $q(n - |M_o|)$ degrees of freedom evaluated at $1/4$. The proof is complete by applying Lemma B.1 to the first term in (14). \square

Proof of Theorem 4.9. Let, $j^* := \arg \min_j \|\widehat{\mathbf{B}}_{\cdot M, j}\|_2$, where $\widehat{\mathbf{B}}_{\cdot M, j}$ is the j th column of the least square coefficient matrix $\widehat{\mathbf{B}}_{\cdot M}$ for model M . Construct the model $M(-1) := M \setminus \{j^*\}$ with $|M| - 1$ covariates from model M . Suppose $\widehat{\mathbf{B}}_{\cdot M(-1)}$ be the least-squared estimator corresponding to the model $M(-1)$. Because \mathbf{B}_{\min} minimizes the objective function corresponding to the h function,

$$\begin{aligned}
 & \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B}_{\min} \mathbf{X} \right) \right\|_{\mathbf{F}}^2 \\
 & \leq \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbb{E}_y \left(\widehat{\mathbf{B}}_{\cdot M(-1)} \mathbf{X}_{M(-1) \cdot} \right) \right) \right\|_{\mathbf{F}}^2.
 \end{aligned}$$

By triangle inequality for the Frobenius norm,

$$\begin{aligned}
 & \mathbb{E} (h_\epsilon(\mathbf{B}_{\cdot M})) \\
 & = \mathbb{P} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbf{B}_{\min} \mathbf{X} \right) \right\|_{\mathbf{F}}^2 \geq \epsilon \right) \\
 & \leq \mathbb{P} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbb{E}_y \left(\widehat{\mathbf{B}}_{\cdot M(-1)} \mathbf{X}_{M(-1) \cdot} \right) \right) \right\|_{\mathbf{F}}^2 \geq \epsilon \right) \\
 & \leq \mathbb{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbb{E}_y \left[\widehat{\mathbf{B}}_{\cdot M} \right] \mathbf{X}_{M \cdot} \right) \right\|_{\mathbf{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 & \quad + \mathbb{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbb{E}_y \left[\widehat{\mathbf{B}}_{\cdot M} \right] \mathbf{X}_{M \cdot} - \mathbb{E}_y \left(\widehat{\mathbf{B}}_{\cdot M(-1)} \mathbf{X}_{M(-1) \cdot} \right) \right) \right\|_{\mathbf{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 & \quad + \mathbb{P} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\mathbf{B}_{\cdot M} \mathbf{X}_{M \cdot} - \widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_{M \cdot} \right) \right\|_{\mathbf{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 & \leq \mathbb{I} \left(\frac{1}{2} \left\| \widehat{\boldsymbol{\Sigma}}_{(M)}^{-1/2} \left(\widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_{M \cdot} - \mathbb{E}_y \left[\widehat{\mathbf{B}}_{\cdot M} \right] \mathbf{X}_{M \cdot} \right) \right\|_{\mathbf{F}}^2 \geq \frac{\epsilon}{9} \right)
 \end{aligned}$$

$$\begin{aligned}
 &+ \mathbb{I} \left(\frac{1}{2} \left\| \widehat{\Sigma}_{(M)}^{-1/2} \mathbf{B}^0_{M_o} \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)}) \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 &+ e^{-\frac{\epsilon(n-|M|)}{36} + \frac{q|M|}{2}} + 2e^{-\frac{(\sqrt{n-|M|}-2\sqrt{q})^2}{8}}, \tag{15}
 \end{aligned}$$

where the first probability is computed by Lemma B.1. Next,

$$\begin{aligned}
 &\mathbb{P}_y \left(\left\{ \left\| \widehat{\Sigma}_{(M)}^{-1/2} (\widehat{\mathbf{B}}_{\cdot M} \mathbf{X}_M - \mathbb{E} [\widehat{\mathbf{B}}_{\cdot M}] \mathbf{X}_M) \right\|_{\mathbb{F}}^2 < \frac{2\epsilon}{9} \right\} \cap \right. \\
 &\quad \left. \left\{ \left\| \widehat{\Sigma}_{(M)}^{-1/2} \mathbf{B}^0_{M_o} \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)}) \right\|_{\mathbb{F}}^2 < \frac{2\epsilon}{9} \right\} \right) \\
 &\geq 1 - \mathbb{P}_y \left(\frac{1}{2} \left\| \widehat{\Sigma}_{(M)}^{-1/2} [\widehat{\mathbf{B}}_{\cdot M} - \mathbb{E} (\widehat{\mathbf{B}}_{\cdot M})] \mathbf{X}_M \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 &\quad - \mathbb{P}_y \left(\frac{1}{2} \left\| \widehat{\Sigma}_{(M)}^{-1/2} \mathbf{B}^0_{M_o} \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)}) \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right). \tag{16}
 \end{aligned}$$

By Condition 4.8 and putting $\tau = 1/2$ in Lemma 4.2,

$$\begin{aligned}
 &\mathbb{P}_y \left(\frac{1}{2} \left\| \widehat{\Sigma}_{(M)}^{-1/2} \mathbf{B}^0_{M_o} \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)}) \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 &\leq \mathbb{P}_y \left(\lambda_{\min} (\widehat{\Sigma}_{(M)}) \leq \frac{9}{2\epsilon} \left\| \mathbf{B}^0_{M_o} \mathbf{X}_{M_o} \cdot (\mathbf{H}_{(M)} - \mathbf{H}_{(M)(-1)}) \right\|_{\mathbb{F}}^2 \right) \\
 &\leq \mathbb{P}_y \left(\lambda_{\min} (\widehat{\Sigma}_{(M)}) \leq \frac{\lambda_v(n-|M|)}{2} \right) \leq e^{-0.04(n-|M|)}. \tag{17}
 \end{aligned}$$

Bounding the first probability in equation (16) by Lemma B.2 and combining equation (15),(16), and (17),

$$\mathbb{P}_y \left[\mathbb{E} (h_{\epsilon}(\mathbf{B}_{\cdot M})) \geq e^{-\frac{\epsilon(n-|M|)}{36} + \frac{q|M|}{2}} + 2e^{-\frac{1}{8}\{\sqrt{n-|M|}-2\sqrt{q}\}^2} \right] \leq V_{7,n,M}$$

where

$$V_{7,n,M} := \exp \left(-\frac{\epsilon(n-|M|)\lambda_v}{36\bar{\lambda}_v} + \frac{q|M|}{2} \right) + 2 \exp(-0.04(n-|M|))$$

Finally,

$$\begin{aligned}
 &\mathbb{P}_y \left[\bigcup_{\substack{M \not\subseteq M_o \\ |M| \leq n^\alpha}} \left\{ \mathbf{Y} : \mathbb{E} (h_{\epsilon}(\mathbf{B}_{\cdot M})) \geq e^{-\frac{\epsilon(n-|M|)}{36} + \frac{q|M|}{2}} + 2e^{-\frac{1}{8}\{\sqrt{n-|M|}-2\sqrt{q}\}^2} \right\} \right] \\
 &\leq \sum_{j=1}^{n^\alpha} \sum_{\substack{M:|M|=j \\ M \not\subseteq M_o}} \left\{ \exp \left(-\frac{\epsilon(n-|M|)\lambda_v}{36\bar{\lambda}_v} + \frac{q|M|}{2} \right) + 2 \exp(-0.04(n-|M|)) \right\}
 \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{j=1}^{n^\alpha} \left\{ \exp\left(-\frac{\epsilon(n-j)\lambda_v}{36\bar{\lambda}_v} + \frac{qj}{2} + j \log p\right) + 2 \exp(-0.04(n-j) + j \log p) \right\} \\
&\leq n^\alpha \left\{ \exp\left(-\frac{\epsilon(n-n^\alpha)\lambda_v}{36\bar{\lambda}_v} + \frac{qn^\alpha}{2} + n^\alpha \log p\right) \right. \\
&\quad \left. + 2 \exp(-0.04(n-n^\alpha) + n^\alpha \log p) \right\}.
\end{aligned}$$

This concludes the proof of the theorem. \square

Proof of Lemma B.1. We define $\tilde{\mathbf{B}}_{\cdot M}$ equal to $\hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} - \hat{\mathbf{B}}_{\cdot M})$ and $\mathbf{\Omega}_{(M)}$ equal to $\mathbf{X}_M \mathbf{X}_M^\top$, observe that,

$$\begin{aligned}
\left\| \hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} \mathbf{X}_M - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M) \right\|_{\mathbb{F}}^2 &= \left\| \hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} - \hat{\mathbf{B}}_{\cdot M}) \mathbf{X}_M \right\|_{\mathbb{F}}^2 \\
&= \text{tr} \left(\tilde{\mathbf{B}}_{\cdot M} \mathbf{\Omega}_{(M)} \tilde{\mathbf{B}}_{\cdot M}^\top \right) \\
&= \left\| \hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} - \hat{\mathbf{B}}_{\cdot M}) \mathbf{\Omega}_{(M)}^{1/2} \right\|_{\mathbb{F}}^2.
\end{aligned}$$

By the property of Matrix-t distribution, (see Theorem 4.3.5, page 137 of [23])

$$\hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} - \hat{\mathbf{B}}_{\cdot M}) \mathbf{\Omega}_{(M)}^{1/2} \sim T_{q,M} (n - |M| - q + 1, \mathbf{0}, \mathbf{I}_q, \mathbf{I}_{|M|}).$$

Additionally, by Theorem 4.2.1 (page 134) of [23],

$$\hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} - \hat{\mathbf{B}}_{\cdot M}) \mathbf{\Omega}_{(M)}^{1/2} \stackrel{d}{=} (\mathbf{W}^{-1/2})^\top \mathbf{Z},$$

where $\mathbf{W} \sim \text{Wishart}_q(n - |M|, \mathbf{I}_q)$ and $\mathbf{Z} \sim \text{Matrix-Normal}(\mathbf{0}, \mathbf{I}_q, \mathbf{I}_{|M|})$. Then,

$$\left\| \hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} \mathbf{X}_M - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M) \right\|_{\mathbb{F}}^2 \leq \lambda_{\min}^{-1}(\mathbf{W}) \text{tr}(\mathbf{Z}\mathbf{Z}^\top),$$

where $\lambda_{\min}(\mathbf{W}) > 0$ with probability 1 as $n - |M| > q$. Since, $\text{tr}(\mathbf{Z}\mathbf{Z}^\top) \sim \chi_{q|M|}^2$,

$$\begin{aligned}
&\mathbb{P} \left(\frac{1}{2} \left\| \hat{\Sigma}_{(M)}^{-1/2} (\mathbf{B}_{\cdot M} \mathbf{X}_M - \hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M) \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right) \\
&\leq \mathbb{P} \left(\frac{\text{tr}(\mathbf{Z}\mathbf{Z}^\top)}{\lambda_{\min}(\mathbf{W})} \geq \frac{2\epsilon}{9} \right) \\
&\leq \mathbb{P} \left(\frac{\text{tr}(\mathbf{Z}\mathbf{Z}^\top)}{\lambda_{\min}(\mathbf{W})} \geq \frac{2\epsilon}{9}, \lambda_{\min}(\mathbf{W}) \geq \frac{(n - |M|)}{4} \right)
\end{aligned}$$

$$\begin{aligned}
 & + \mathbb{P} \left(\frac{\text{tr}(\mathbf{Z}\mathbf{Z}^\top)}{\lambda_{\min}(\mathbf{W})} \geq \frac{2\epsilon}{9}, \lambda_{\min}(\mathbf{W}) < \frac{(n-|M|)}{4} \right) \\
 & \leq \mathbb{P} \left(\text{tr}(\mathbf{Z}\mathbf{Z}^\top) \geq \frac{\epsilon(n-|M|)}{18} \right) + \mathbb{P} \left(\lambda_{\min}(\mathbf{W}) < \frac{1}{4}(n-|M|) \right) \\
 & \leq \mathbb{P} \left(\chi_{q|M|}^2 \geq \frac{\epsilon(n-|M|)}{18} \right) + \mathbb{P} \left(\lambda_{\min}(\mathbf{W}) < \frac{1}{4}(n-|M|) \right) \\
 & \leq \exp \left(-\frac{\epsilon(n-|M|)}{72} + \frac{q|M|}{2} \right) + 2 \exp \left(-\frac{(\sqrt{n-|M|} - 2\sqrt{q})^2}{8} \right),
 \end{aligned}$$

where the first probability in the second to last line is obtained by using the Chernoff's bound for χ^2 distribution and the second probability by substituting $t = (1/2)\sqrt{n-|M|} - \sqrt{q} > 0$ (by the condition on $|M|$ in the statement of the lemma) in the corollary 5.35 (page 21) of [49] and noting that any central Wishart $_q(n-|M|, \mathbf{I}_q)$ matrix is identically distributed as $\tilde{\mathbf{Z}}^\top \tilde{\mathbf{Z}}$ where $\tilde{\mathbf{Z}} \sim \text{Matrix-Normal}_{n-|M|,q}(0, \mathbf{I}_{n-|M|}, \mathbf{I}_q)$, provided $n-|M| > q$ which is again true by the specified condition in the lemma. This completes the proof. \square

Proof of Lemma B.2. Defining, $\mathbf{\Omega}_{(M)} = \mathbf{X}_M \cdot \mathbf{X}_M^\top$, the quadratic form can be alternatively written as

$$\begin{aligned}
 & \left\| \hat{\mathbf{\Sigma}}_{(M)}^{-1/2} \left[\hat{\mathbf{B}}_{\cdot M} \mathbf{X}_M \cdot - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \mathbf{X}_M \cdot \right] \right\|_{\mathbb{F}}^2 \\
 & = \text{tr} \left[\hat{\mathbf{\Sigma}}_{(M)}^{-1} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right) \mathbf{\Omega}_{(M)} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right)^\top \right] \\
 & \leq \lambda_{\min}^{-1} \left(\hat{\mathbf{\Sigma}}_{(M)} \right) \bar{\lambda}_v \left\| \mathbf{A}_{(M_o)}^{0^{-1}} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right) \mathbf{\Omega}_{(M)}^{1/2} \right\|_{\mathbb{F}}^2.
 \end{aligned}$$

Since, $\mathbf{A}_{(M_o)}^{0^{-1}} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right) \mathbf{\Omega}_{(M)}^{1/2} \sim \text{Matrix-Normal}_{q,|M|}(\mathbf{0}, \mathbf{I}_q, \mathbf{I}_{|M|})$,

$$\left\| \mathbf{A}_{(M_o)}^{0^{-1}} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right) \mathbf{\Omega}_{(M)}^{1/2} \right\|_{\mathbb{F}}^2 \sim \chi_{q|M|}^2.$$

Choosing $\tau = 1/2$ in Lemma 4.2,

$$\mathbb{P}_y \left(\lambda_{\min} \left(\hat{\mathbf{\Sigma}}_{(M)} \right) < \lambda_v(n-|M|)/2 \right) \leq \exp(-0.4(n-|M|)).$$

Finally,

$$\begin{aligned}
 & \mathbb{P}_y \left(\frac{1}{2} \left\| \hat{\mathbf{A}}_{(M)}^{-1} \left[\hat{\mathbf{B}}_{\cdot M} - \mathbb{E} \left(\hat{\mathbf{B}}_{\cdot M} \right) \right] \mathbf{X}_M \cdot \right\|_{\mathbb{F}}^2 \geq \frac{\epsilon}{9} \right) \\
 & \leq \mathbb{P}_y \left(\lambda_{\min}^{-1} \left(\hat{\mathbf{\Sigma}}_{(M)} \right) \bar{\lambda}_v \left\| \mathbf{V}_{(M_o)}^{0^{-1/2}} \left(\hat{\mathbf{B}}_{\cdot M} - \mathbb{E}_y \left(\hat{\mathbf{B}}_{\cdot M} \right) \right) \mathbf{\Omega}_{(M)}^{1/2} \right\|_{\mathbb{F}}^2 \geq \frac{2\epsilon}{9} \right)
 \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{P}_y \left(\frac{\bar{\lambda}_v \chi_{q|M|}^2}{\lambda_{\min}(\widehat{\Sigma}_{(M)})} \geq \frac{2\epsilon}{9} \right) \\
&\leq \mathbb{P}_y \left(\frac{\bar{\lambda}_v \chi_{q|M|}^2}{\lambda_{\min}(\widehat{\Sigma}_{(M)})} \geq \frac{2\epsilon}{9}, \lambda_{\min}(\widehat{\Sigma}_{(M)}) > \frac{\lambda_v(n-|M|)}{2} \right) \\
&\quad + \mathbb{P}_y \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) \leq \frac{\lambda_v(n-|M|)}{2} \right) \\
&\leq \mathbb{P}_y \left(\chi_{q|M|}^2 \geq \frac{\epsilon(n-|M|)\lambda_v}{9\bar{\lambda}_v} \right) + \mathbb{P}_y \left(\lambda_{\min}(\widehat{\Sigma}_{(M)}) \leq \frac{\lambda_v(n-|M|)}{2} \right) \\
&\leq \exp \left(-\frac{\epsilon(n-|M|)\lambda_v}{36\bar{\lambda}_v} + \frac{q|M|}{2} \right) + \exp(-0.04(n-|M|)),
\end{aligned}$$

where the first quantity in the last line is obtained by Chernoff's bound for chi-square distribution. This completes the proof of Lemma B.2. \square

Proof of Theorem 4.10. The statement of the theorem is equivalent to showing that,

$$\sum_{\substack{M: |M| \leq n^\alpha \\ M \neq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} \xrightarrow{\mathbb{P}_y} 0, \tag{18}$$

as $n \rightarrow \infty$ or $n, p \rightarrow \infty$. To show this, observe that the ratio in sum has the form,

$$\begin{aligned}
\frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} &= \pi^{\frac{q(|M|-|M_o|)}{2}} \frac{\Gamma_q \left(\frac{n-|M|}{2} \right)}{\Gamma_q \left(\frac{n-|M_o|}{2} \right)} \left[\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \widehat{\Sigma}_{(M)}} \right]^{\frac{n-|M|-q}{2}} \times \\
&\quad \left(\det \widehat{\Sigma}_{(M_o)} \right)^{\frac{|M|-|M_o|}{2}} \frac{\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M))}{\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M_o))}.
\end{aligned}$$

The multivariate gamma function is defined as the product of univariate gamma functions, and so, using the gamma function inequalities in [25], the ratio of the multivariate gamma functions is bounded by,

$$\begin{aligned}
\frac{\Gamma_q \left(\frac{n-|M|}{2} \right)}{\Gamma_q \left(\frac{n-|M_o|}{2} \right)} &= \prod_{j=1}^q \Gamma \left(\frac{n-|M|-j+1}{2} \right) \left[\Gamma \left(\frac{n-|M_o|-j+1}{2} \right) \right]^{-1} \\
&\leq \begin{cases} \prod_{j=1}^q \left(\frac{n-|M_o|-j+1}{2} \right) \left(\frac{n-|M|-j+1}{2} \right)^{\frac{|M_o|-|M|}{2}-1} & \text{if } |M_o| \geq |M| \\ \prod_{j=1}^q \left(\frac{n-|M|-j+1}{2} - 1 \right)^{\frac{|M_o|-|M|}{2}} & \text{otherwise} \end{cases}
\end{aligned}$$

Moreover, for the true model M_o , because $\widehat{\Sigma}_{(M_o)} \sim \text{Wishart}_q(n - |M_o|, \mathbf{V}_{(M_o)}^0)$,

$$\frac{\det \widehat{\Sigma}_{(M_o)}}{\det \mathbf{V}_{(M_o)}^0} \sim \prod_{j=1}^q \chi_{n-|M_o|-j+1}^2,$$

by Theorem 3.2.15 of [35]. Applying the Chernoff bound for the chi-square distribution and using the sub-additivity property of the probability measure,

$$\begin{aligned} & \mathbb{P}_y \left(\det \widehat{\Sigma}_{(M_o)} > \prod_{j=1}^q 3\bar{\lambda}_v (n - |M_o| - j + 1) \right) \\ & \leq q \exp \left(-\frac{n - |M_o| - q + 1}{4} \right) := V_{5,n}. \end{aligned}$$

Expanding the quantity in (18),

$$\begin{aligned} \sum_{\substack{M:|M|\leq n^\alpha \\ M \neq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} &= \sum_{M: M \subsetneq M_o} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} + \sum_{\substack{M:|M|\leq n^\alpha \\ M \not\subseteq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} \\ &= \underbrace{\sum_{j=1}^{|M_o|} \sum_{\substack{M:|M|=j \\ M \subsetneq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})}}_{:= T_1} + \underbrace{\sum_{j=1}^{n^\alpha} \sum_{\substack{M:|M|=j \\ M \not\subseteq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})}}_{:= T_2}. \end{aligned} \tag{19}$$

Denote the two terms on the right side as T_1 and T_2 , respectively. First consider T_1 . By Theorem 4.7, with probability exceeding $1 - V_{3,n}$, $\mathbb{E}(h_\epsilon(\mathbf{B}_{M_o}))$ is bounded from below by $1 - g_n(M_o, \epsilon)$ with,

$$g_n(M_o, \epsilon) := \exp \left(-\frac{\epsilon(n - |M_o|)}{36} + \frac{q|M_o|}{2} \right) + 2 \exp \left(-\frac{1}{8} \left\{ \sqrt{n - |M_o|} - 2\sqrt{q} \right\}^2 \right).$$

Since the quantity $g_n(M_o, \epsilon)$ vanishes as $n \rightarrow \infty$, for sufficiently large n , $g_n(M_o, \epsilon)$ less than K for some $K \in (0, 1)$. Bounding the ratio of the determinants of the residual matrices by Case 1 of Lemma 4.5, choosing n large enough so that $\mathbb{E}(h_\epsilon(\mathbf{B}_{M_o})) > 1 - K$, and bounding the $\mathbb{E}(h_\epsilon(\mathbf{B}_M))$ for any model $M \not\subseteq M_o$ by 1,

$$\begin{aligned} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} &\leq \frac{e^{-qn^\alpha \log|M_o|}}{1 - K} \prod_{j=1}^q \frac{\left(\frac{n - |M_o| - j + 1}{2} \right) (3\pi\bar{\lambda}_v (n - |M_o| - j + 1))^{\frac{|M| - |M_o|}{2}}}{\left(\frac{n - |M| - j + 1}{2} \right)^{1 - \frac{|M_o| - |M|}{2}}} \\ &\leq \frac{e^{-qn^\alpha \log|M_o|}}{1 - K} \prod_{j=1}^q \left[\frac{n - |M| - j + 1}{6\pi\bar{\lambda}_v (n - |M_o| - j + 1)} \right]^{\frac{|M_o| - |M|}{2}} \end{aligned}$$

$$\begin{aligned} &\leq \frac{e^{-qn^\alpha \log|M_o|}}{1-K} \prod_{j=1}^q \left[\frac{n}{6\pi\bar{\lambda}_v(n-|M_o|-j+1)} \right]^{\frac{|M_o|-|M|}{2}} \\ &\leq \frac{1}{1-K} e^{-qn^\alpha \log|M_o|-q \log(\bar{\lambda}_v)(|M_o|-|M|)/2}, \end{aligned}$$

where the last inequality holds for $n > \frac{|M_o|+q}{1-1/6\pi}$. This implies that with probability exceeding $1 - V_{1,n} - V_{3,n} - V_{5,n}$,

$$\begin{aligned} T_1 &= \sum_{j=1}^{|M_o|} \sum_{\substack{M:|M|=j \\ M \not\subseteq M_o}} \frac{r_\epsilon(M|\mathbf{Y})}{r_\epsilon(M_o|\mathbf{Y})} \leq \sum_{j=1}^{|M_o|} \binom{|M_o|}{j} \max_{M \subseteq M_o; |M|=j} \frac{r_\epsilon(M|\mathbf{Y})}{r_\epsilon(M_o|\mathbf{Y})} \\ &\leq \sum_{j=1}^{|M_o|} (1-K)^{-1} e^{-qn^\alpha \log|M_o|-q \log(\bar{\lambda}_v)(|M_o|-j)/2+j \log|M_o|} \\ &\leq \frac{1}{K-1} |M_o| e^{-qn^\alpha \log|M_o|+|M_o| \log|M_o|-qI(\bar{\lambda}_v < 1) \log(\bar{\lambda}_v)|M_o|/2}. \end{aligned}$$

By Condition 4.3, $1 - V_{1,n} - V_{3,n} - V_{5,n} \rightarrow 1$ as $n \rightarrow \infty$ or $n, p \rightarrow \infty$, and so $T_1 \rightarrow 0$ in probability as $n \rightarrow \infty$ or $n, p \rightarrow \infty$.

Next, consider T_2 . Note that for models M such that $M \not\subseteq M_o$ and $|M| \leq n^\alpha$, By Theorem 4.9, for large n , with probability exceeding $1 - V_{4,n}$, $\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M)) \leq \tilde{g}_n(\epsilon, M)$ where,

$$\tilde{g}_n(\epsilon, M) := \exp\left(-\frac{\epsilon(n-|M|)}{36} + \frac{q|M|}{2}\right) + 2 \exp\left(-\frac{1}{8} \left\{ \sqrt{n-|M|} - 2\sqrt{q} \right\}^2\right).$$

Bounding the ratio of the determinants of the residual matrices as in Case 2 of Theorem 4.5, and choosing n large enough so that $\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M_o)) > 1 - K$ with high probability, for all $M \not\subseteq M_o$ with $|M| \leq n^\alpha$,

$$\frac{r_\epsilon(M|\mathbf{Y})}{r_\epsilon(M_o|\mathbf{Y})} \leq \begin{cases} \left[\frac{n}{6\pi\bar{\lambda}_v(n-|M_o|-q)} \right]^{\frac{q(|M_o|-|M|)}{2}} \frac{\tilde{g}_n(\epsilon, M)}{1-K} e^{q(n^\alpha \log(n-|M_o|)+|M| \log p)} & \text{if } |M| \leq |M_o| \\ \left(\frac{6\pi\bar{\lambda}_v(n-|M_o|)}{n-|M|-q-1} \right)^{q \left(\frac{|M|-|M_o|}{2} \right)} \frac{\tilde{g}_n(\epsilon, M)}{1-K} e^{q(n^\alpha \log(n-|M_o|)+|M| \log p)} & \text{if otherwise.} \end{cases}$$

$$T_2 = \underbrace{\sum_{j=1}^{|M_o|} \sum_{\substack{M:|M|=j \\ M \not\subseteq M_o}} \frac{r_\epsilon(M|\mathbf{Y})}{r_\epsilon(M_o|\mathbf{Y})}}_{:= T_{21}} + \underbrace{\sum_{j=|M_o|+1}^{n^\alpha} \sum_{\substack{M:|M|=j \\ M \not\subseteq M_o}} \frac{r_\epsilon(M|\mathbf{Y})}{r_\epsilon(M_o|\mathbf{Y})}}_{:= T_{22}}$$

Consider T_{21} and T_{22} separately. For sufficiently large n , using $\tilde{g}_n(\epsilon, M)$ to bound

$\mathbb{E}(h_\epsilon(\mathbf{B}_\cdot M))$, with probability exceeding $1 - V_{2,n} - V_{3,n} - V_{4,n} - V_{5,n}$,

$$\begin{aligned}
T_{21} &\leq \sum_{j=1}^{|M_o|} \binom{p}{j} \max_{M \not\subseteq M_o; |M|=j} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} \\
&\leq \frac{2}{1-K} \sum_{j=1}^{|M_o|} e^{-\frac{n-j}{8} - q \log(\bar{\lambda}_v)(|M_o|-j)/2 + qn^\alpha \log(n-|M_o|) + (q+1)j \log p} \\
&\quad + \frac{1}{1-K} \sum_{j=1}^{|M_o|} e^{-\frac{\epsilon(n-j)}{36} + \frac{qj}{2} - q \log(\bar{\lambda}_v)(|M_o|-j)/2 + qn^\alpha \log(n-|M_o|) + (q+1)j \log p} \\
&\leq \frac{|M_o|}{1-K} \left\{ 2e^{-\frac{n-|M_o|}{8} - \frac{q|M_o|}{2} \mathbf{I}(\bar{\lambda}_v < 1) \log(\bar{\lambda}_v) + qn^\alpha \log(n-|M_o|) + (q+1)|M_o| \log p} \right. \\
&\quad \left. + e^{-\frac{\epsilon(n-|M_o|)}{36} + qn^\alpha \log(n-|M_o|) + (q+1)|M_o| \log p + \frac{q|M_o|}{2} \{1 - \mathbf{I}(\bar{\lambda}_v < 1) \log(\bar{\lambda}_v)\}} \right\}.
\end{aligned}$$

For T_{22} , because $(n - |M_o|)/(n - |M| - q - 1)$ converges to 1 for all M with $|M| \leq n^\alpha$, we can choose n large enough so that $(n - |M_o|)/(n - |M| - q - 1) < 2$. Then, again using the bound $\tilde{g}_n(\epsilon, M)$, we obtain that for all $M \not\subseteq M_o$ such that $|M_o| \leq |M| \leq n^\alpha$, with probability exceeding $1 - V_{2,n} - V_{3,n} - V_{4,n} - V_{5,n}$,

$$\begin{aligned}
T_{22} &:= \sum_{j=|M_o|+1}^{n^\alpha} \sum_{\substack{M: |M|=j \\ M \not\subseteq M_o}} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} \\
&\leq \sum_{j=1}^{n^\alpha} \binom{p}{j} \max_{M \not\subseteq M_o; |M|=j} \frac{r_\epsilon(M | \mathbf{Y})}{r_\epsilon(M_o | \mathbf{Y})} \\
&\leq \frac{2}{1-K} \sum_{j=1}^{n^\alpha} e^{-\frac{n-j}{8} + \frac{qj}{2} \log(12\pi\bar{\lambda}_v) + qn^\alpha \log(n-|M_o|) + (q+1)j \log p} \\
&\quad + \frac{1}{1-K} \sum_{j=1}^{n^\alpha} e^{-\frac{\epsilon(n-j)}{36} + \frac{qj}{2} (1 + \log(12\pi\bar{\lambda}_v)) + qn^\alpha \log(n-|M_o|) + (q+1)j \log p} \\
&\leq \frac{2n^\alpha}{1-K} \exp\left(-\frac{n-n^\alpha}{8} + n^\alpha \{q \log(n-|M_o|) + (q+1) \log p\}\right) \\
&\quad \quad \quad + \frac{q}{2} \log(12\pi\bar{\lambda}_v) \left\{ n^\alpha \mathbf{I}(\bar{\lambda}_v > \frac{1}{12\pi}) + 1 \right\} \\
&\quad + \frac{n^\alpha}{1-K} \exp\left(-\frac{\epsilon(n-n^\alpha)}{36} + qn^\alpha \log(n-|M_o|) + (q+1)n^\alpha \log p\right) \\
&\quad \quad \quad + \frac{q}{2} (1 + \log(12\pi\bar{\lambda}_v)) \left\{ n^\alpha \mathbf{I}(\bar{\lambda}_v > 1/12\pi e) + 1 \right\}.
\end{aligned}$$

Lastly, by Condition 4.3 since $1 - V_{2,n} - V_{3,n} - V_{4,n} - V_{5,n} \rightarrow 1$ as $n \rightarrow \infty$ or $n, p \rightarrow \infty$, $T_2 = T_{21} + T_{22} \rightarrow 0$ in probability as $n \rightarrow \infty$ or $n, p \rightarrow \infty$. This completes the proof of Theorem 4.10. \square

Acknowledgments

The authors would like to thank the associate editor and two reviewers for their constructive comments which led to a significantly improved version of the manuscript. The authors would also like to thank Ms. Sukanya Bhattacharyya for providing additional computational resources without which the extensive numerical studies presented in the paper would not have been possible.

Research reported in this publication was supported by the National Heart, Lung, and Blood Institute of the National Institutes of Health under Award Number R56HL155373. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- [1] BAI, R. and GHOSH, M. (2018). High-dimensional multivariate posterior consistency under global-local shrinkage priors. *Journal of Multivariate Analysis* **167** 157–170. [MR3830639](#)
- [2] BAI, R. and GHOSH, M. (2018). MBSP: Multivariate Bayesian Model with Shrinkage Priors R package version 1.0.
- [3] BAI, R., MORAN, G. E., ANTONELLI, J. L., CHEN, Y. and BOLAND, M. R. (2020). Spike-and-slab group lassos for grouped regression and sparse generalized additive models. *Journal of the American Statistical Association* 1–14. [MR4399078](#)
- [4] BELLEC, P. C. and ROMON, G. (2021). Chi-square and normal inference in high-dimensional multi-task regression. *arXiv preprint arXiv:2107.07828*.
- [5] BERTRAND, Q., MASSIAS, M., GRAMFORT, A. and SALMON, J. (2019). Handling correlated and repeated measurements with the smoothed multivariate square-root Lasso. *arXiv preprint arXiv:1902.02509*.
- [6] BERTSIMAS, D., KING, A. and MAZUMDER, R. (2016). Best subset selection via a modern optimization lens. *The Annals of Statistics* **44** 813–852. [MR3476618](#)
- [7] BOULESTEIX, A.-L. and STRIMMER, K. (2005). Predicting transcription factor activities from combined analysis of microarray and ChIP data: a partial least squares approach. *Theoretical Biology and Medical Modelling* **2** 1–12.
- [8] BREIMAN, L. and FRIEDMAN, J. H. (1997). Predicting multivariate responses in multiple linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **59** 3–54. [MR1436554](#)
- [9] BROWN, P. J., VANNUCCI, M. and FEARN, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60** 627–641. [MR1626005](#)
- [10] BROWN, P. J., VANNUCCI, M. and FEARN, T. (2002). Bayes model averaging with selection of regressors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64** 519–536. [MR1924304](#)

- [11] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media. [MR2807761](#)
- [12] CHEN, K. (2019). rrpck: Reduced-Rank Regression R package version 0.1-11.
- [13] CHEN, L. and HUANG, J. Z. (2012). Sparse reduced-rank regression for simultaneous dimension reduction and variable selection. *Journal of the American Statistical Association* **107** 1533–1545. [MR3036414](#)
- [14] CHEVALIER, J.-A., GRAMFORT, A., SALMON, J. and THIRION, B. (2020). Statistical control for spatio-temporal MEG/EEG source imaging with desparsified multi-task Lasso. *arXiv preprint arXiv:2009.14310*. [MR3829759](#)
- [15] CHUN, H. and KELEŞ, S. (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72** 3–25. [MR2751241](#)
- [16] CHUNG, D., CHUN, H. and KELES, S. (2019). spls: Sparse Partial Least Squares (SPLS) Regression and Classification R package version 2.2-3.
- [17] DESHPANDE, S. K., ROČKOVÁ, V. and GEORGE, E. I. (2019). Simultaneous variable and covariance selection with the multivariate spike-and-slab lasso. *Journal of Computational and Graphical Statistics* **28** 921–931. [MR4045858](#)
- [18] FRANK, L. E. and FRIEDMAN, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35** 109–135.
- [19] FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software* **33** 1–22.
- [20] FUJIKOSHI, Y. and SATOH, K. (1997). Modified AIC and Cp in multivariate linear regression. *Biometrika* **84** 707–716. [MR1603952](#)
- [21] GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. and RUBIN, D. B. (2013). *Bayesian data analysis (3rd ed.)*. Chapman and Hall/CRC. [MR3235677](#)
- [22] GUPTA, A. K. and NAGAR, D. K. (2018). *Matrix variate distributions* **104**. CRC Press. [MR1738933](#)
- [23] GUPTA, A. K. and NAGAR, D. K. (2018). *Matrix variate distributions* **104**. CRC Press. [MR1738933](#)
- [24] HANNIG, J., IYER, H., LAI, R. C. and LEE, T. C. (2016). Generalized fiducial inference: A review and new results. *Journal of the American Statistical Association* **111** 1346–1361. [MR3561954](#)
- [25] JAMESON, G. (2013). Inequalities for gamma function ratios. *The American Mathematical Monthly* **120** 936–940. [MR3139572](#)
- [26] LAHIRI, S. N. (2021). Necessary and sufficient conditions for variable selection consistency of the LASSO in high dimensions. *The Annals of Statistics* **49** 820–844. [MR4255109](#)
- [27] LEE, T. I., RINALDI, N. J., ROBERT, F., ODOM, D. T., BAR-JOSEPH, Z., GERBER, G. K., HANNETT, N. M., HARBISON, C. T., THOMPSON, C. M.,

- SIMON, I. et al. (2002). Transcriptional regulatory networks in *Saccharomyces cerevisiae*. *science* **298** 799–804.
- [28] LEE, W. and LIU, Y. (2012). Simultaneous multiple response regression and inverse covariance matrix estimation via penalized Gaussian maximum likelihood. *Journal of Multivariate Analysis* **111** 241–255. [MR2944419](#)
- [29] LI, Y., NAN, B. and ZHU, J. (2015). Multivariate sparse group lasso for the multivariate multiple linear regression with an arbitrary group structure. *Biometrics* **71** 354–363. [MR3366240](#)
- [30] LI, Y., NAN, B. and ZHU, J. (2016). MSGLasso: Multivariate Sparse Group Lasso for the Multivariate Multiple Linear Regression with an Arbitrary Group Structure R package version 2.1. [MR3366240](#)
- [31] LIQUET, B., MENGERSEN, K., PETTITT, A., SUTTON, M. et al. (2017). Bayesian variable selection regression of multivariate responses for group data. *Bayesian Analysis* **12** 1039–1067. [MR3724978](#)
- [32] LIQUET, B. and SUTTON, M. (2017). MBSGS: Multivariate Bayesian Sparse Group Selection with Spike and Slab R package version 1.1.0.
- [33] MASSIAS, M., FERCOQ, O., GRAMFORT, A. and SALMON, J. (2018). Generalized concomitant multi-task lasso for sparse multimodal regression. In *International Conference on Artificial Intelligence and Statistics* 998–1007. PMLR.
- [34] MOLSTAD, A. J. (2022). New Insights for the Multivariate Square-Root Lasso. *Journal of Machine Learning Research* **23** 1–52. [MR4576651](#)
- [35] MUIRHEAD, R. J. (2009). *Aspects of multivariate statistical theory* **197**. John Wiley & Sons. [MR0652932](#)
- [36] NARISSETTY, N. N. and HE, X. (2014). Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics* **42** 789–817. [MR3210987](#)
- [37] PENG, J., ZHU, J., BERGAMASCHI, A., HAN, W., NOH, D.-Y., POLLACK, J. R. and WANG, P. (2010). Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *The Annals of Applied Statistics* **4** 53. [MR2758084](#)
- [38] PHILLIPS, T. (2008). Regulation of transcription and gene expression in eukaryotes. *Nature Education* **1** 199.
- [39] ROTHMAN, A. J., LEVINA, E. and ZHU, J. (2010). Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics* **19** 947–962. [MR2791263](#)
- [40] SIMILA, T. and TIKKA, J. (2006). Common subset selection of inputs in multiresponse regression. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings* 1908–1915. IEEE.
- [41] SONDEREGGER, D. L. and HANNIG, J. (2014). Fiducial theory for free-knot splines. In *Contemporary Developments in Statistical Theory* 155–189. Springer. [MR3149921](#)
- [42] SPARKS, R., COUTSOURIDES, D. and TROSKIE, L. (1983). The multivariate CP. *Communications in Statistics-Theory and Methods* **12** 1775–1793. [MR0704853](#)
- [43] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso.

- Journal of the Royal Statistical Society: Series B (Methodological)* **58** 267–288. [MR1379242](#)
- [44] TU, S. (2020). On the Smallest Singular Value of Non-Centered Gaussian Designs. https://stephentu.github.io/writeups/non_centered_gaussian.pdf.
- [45] TURLACH, B. A., VENABLES, W. N. and WRIGHT, S. J. (2005). Simultaneous variable selection. *Technometrics* **47** 349–363. [MR2164706](#)
- [46] VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42** 1166–1202. [MR3224285](#)
- [47] VAN DE GEER, S. and STUCKY, B. (2016). χ^2 -confidence sets in high-dimensional regression. In *Statistical analysis for high-dimensional data* 279–306. Springer. [MR3616273](#)
- [48] VELU, R. and REINSEL, G. C. (2013). *Multivariate reduced-rank regression: theory and applications* **136**. Springer Science & Business Media. [MR1719704](#)
- [49] VERSHYNIN, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*. [MR2963170](#)
- [50] VOUNOU, M., NICHOLS, T. E., MONTANA, G. and INITIATIVE, A. D. N. (2010). Discovering genetic associations with high-dimensional neuroimaging phenotypes: a sparse reduced-rank regression approach. *Neuroimage* **53** 1147–1159.
- [51] WANG, G., WANG, F., HUANG, Q., LI, Y., LIU, Y. and WANG, Y. (2015). Understanding transcription factor regulation by integrating gene expression and dnase i hypersensitive sites. *BioMed research international* **2015**.
- [52] WANG, L., CHEN, G. and LI, H. (2007). Group SCAD regression analysis for microarray time course gene expression data. *Bioinformatics* **23** 1486–1494.
- [53] WILLIAMS, J. P. and HANNIG, J. (2019). Nonpenalized variable selection in high-dimensional linear model settings via generalized fiducial inference. *The Annals of Statistics* **47** 1723–1753. [MR3911128](#)
- [54] WILLIAMS, J. P., XIE, Y. and HANNIG, J. (2019). The EAS approach for graphical selection consistency in vector autoregression models. *arXiv preprint arXiv:1906.04812*. [MR4595246](#)
- [55] WILMS, I. and CROUX, C. (2018). An algorithm for the multivariate group lasso with covariance estimation. *Journal of Applied Statistics* **45** 668–681. [MR3750468](#)
- [56] ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *The Journal of Machine Learning Research* **7** 2541–2563. [MR2274449](#)