

Factor-Adjusted Regularized Model Selection *

Jianqing Fan

Department of ORFE, Princeton University

Yuan Ke

Department of Statistics, University of Georgia

Kaizheng Wang

Department of ORFE, Princeton University

Abstract

This paper studies model selection consistency for high dimensional sparse regression when data exhibits both cross-sectional and serial dependency. Most commonly-used model selection methods fail to consistently recover the true model when the covariates are highly correlated. Motivated by econometric and financial studies, we consider the case where covariate dependence can be reduced through factor model, and propose a consistency strategy named Factor-Adjusted Regularized Model Selection (FarmSelect). By learning the latent factors and idiosyncratic components and using both of them as predictors, FarmSelect transforms the problem from model selection with highly correlated covariates to that with weakly correlated ones via lifting. Model selection consistency as well as optimal rates of convergence are obtained under mild conditions. Numerical studies demonstrate the nice finite sample performance in terms of both model selection and out-of-sample prediction. Moreover, our method is flexible in a sense that it pays no price for weakly correlated and uncorrelated cases. Our method is applicable to a wide range of high dimensional sparse regression problems. An R-package *FarmSelect* is also provided for implementation.

Key words: High dimension; Model selection consistency; Correlated covariates; Factor model; Regularized M -estimator; Time series.

*Corresponding author. Address: 310 Herty Drive University of Georgia, Athens, GA 30602, USA. Phone: 609-955-8395. Fax: 706-542-3391. Email: yuan.ke@uga.edu. The research was supported in part by the NSF grants DMS-1662139 and DMS-1712591, by the ONR grant N00014-19-1-2120, and by the NIH grant 2R01-GM072611-13.

1 Introduction

With the development of data collection and storage technologies, high dimensional time series characterize many contemporary research problems in economics, finance, genomics, statistics, machine learning and so on. Specifying an appropriate yet parsimonious model has become a key topic in high dimensional time series analysis. Parsimonious models are preferable due to their simplicity and interpretability. In classic econometric studies, extensive efforts have been made to identify the correct orders of time series models, see [Akaike \(1973\)](#), [Schwarz \(1978\)](#), [Tsay and Tiao \(1985\)](#), [Choi \(1992\)](#) and [Tiao and Tsay \(1989\)](#) among others. In addition, removing redundant coefficients can improve the prediction accuracy of time series. Professor George C. Tiao and his co-authors, among others, have contributed to this area by a series of pioneering works ([Box and Tiao, 1976](#); [Liu et al., 1992](#); [Montgomery et al., 1998](#)).

Over the past two decades, many model selection methods have been developed. A major part of them are based on the regularized M -estimation approach including the LASSO ([Tibshirani, 1996](#)), the SCAD ([Fan and Li, 2001](#)), the elastic net ([Zou and Hastie, 2005](#)), and the Dantzig selector ([Candes and Tao, 2007](#)), among others. These methods have attracted a large amount of theoretical and algorithmic studies. See [Donoho and Elad \(2003\)](#), [Fan and Peng \(2004\)](#), [Efron et al. \(2004\)](#), [Meinshausen and Bühlmann \(2006\)](#), [Zhao and Yu \(2006\)](#), [Fan and Lv \(2008\)](#), [Zou and Li \(2008\)](#), [Bickel et al. \(2009\)](#), [Wainwright \(2009\)](#), [Zhang \(2010\)](#), and references therein. However, most existing model selection schemes are not tailored for economic and finance applications as they assume covariates are cross-sectionally weakly correlated and serially independent. These conditions are easily violated in economic and financial datasets. For example, economics studies (e.g. [Stock and Watson, 2002](#); [Bai and Ng, 2002](#)) show that there exist strong co-movements among a large pool of macroeconomic variables. A stylized feature of the stock return data is cross-sectionally correlated among the stock returns. Furthermore, even if the weakly correlated assumption holds, one may still observe strong spurious correlations in a high dimensional sample.

To illustrate how cross-sectional correlations influence the model selection result, we consider a toy example of LASSO with an equally correlated design. Consider a sparse linear model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}$. We choose sample size $n = 100$, dimensionality $p = 200$, $\boldsymbol{\beta}^* = (\beta_1, \dots, \beta_{10}, \mathbf{0}_{(p-10)}^T)^T$, and $\boldsymbol{\varepsilon} \sim N(\mathbf{0}_n, 0.3\mathbf{I}_n)$. The nonzero coefficients $\beta_1, \dots, \beta_{10}$ are drawn from i.i.d. Uniform $[2, 5]$. The covariates $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p)^T$ are drawn from the normal distribution $N(\mathbf{0}_p, \boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma}$ is a correlation matrix with all off-diagonal elements ρ for some $\rho \in [0, 1]$. Let ρ increase from 0 to 0.95 by a step size 0.05. For each given ρ , we simulate 200 replications and calculate the average




Figure 1: LASSO model selection results with respect to the correlations

model size selected by LASSO, the average model size when the first false discovery (\mathbf{x}_j , $j > 10$) enters the solution path and the model selection consistency rate. As is shown in Figure 1, the correlation influences the model selection results in the following three aspects: (i) selected model size, (ii) early selection of false variables, (iii) model selection consistency rates. Therefore, when the covariates are highly correlated, there is little hope to exactly recover the active set from the solution path of LASSO. As to be shown later, the correlation has similar adverse impacts on other model selection methods (e.g. SCAD and elastic net).

To overcome the the aforementioned problems caused by the cross-sectional correlation, this paper proposes a consistent strategy named Factor-Adjusted Regularized Model Selection (Farm-Select) for the case where covariates can be decorrelated via a few pervasive latent factors. More precisely, let x_{tj} be the t th ($t = 1, \dots, n$) observation of the j th ($j = 1, \dots, p$) covariate, and assume that $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})^T$ follows an approximate factor model

$$\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t, \quad (1.1)$$

where \mathbf{f}_t is a $K \times 1$ vector of latent factors, \mathbf{B} is a $p \times K$ matrix of factor loadings, and \mathbf{u}_t is a $p \times 1$ vector of idiosyncratic components that are uncorrelated with \mathbf{f}_t . The strategy of FarmSelect is to first learn the parameters in approximate factor model (1.1) for the covariates $\{\mathbf{x}_t\}_{t=1}^n$. Denote by $\hat{\mathbf{f}}_t$ and $\hat{\mathbf{B}}$ the obtained estimators of the factors and loadings respectively. Then by identifying the

highly correlated low rank part by $\widehat{\mathbf{B}}\widehat{\mathbf{f}}_t$, we transform the problem from model selection with highly correlated covariates in \mathbf{x}_t to that with weakly correlated or uncorrelated idiosyncratic components $\widehat{\mathbf{u}}_t := \mathbf{x}_t - \widehat{\mathbf{B}}\widehat{\mathbf{f}}_t$ and $\widehat{\mathbf{f}}_t$. This lifting step makes covariates weakly correlated. The second step amounts to solving a regularized profile likelihood problem. We study FarmSelect in details by providing theoretical guarantees that FarmSelect can achieve model selection consistency as well as estimation consistency under mild conditions. Unlike traditional studies in model selection where the samples are assumed to be i.i.d., serial dependency is allowed and thus our theories apply to time series data. Moreover, both theoretical and numerical studies show the flexibility of FarmSelect in a sense that it pays no price for weakly correlated cases. This property makes FarmSelect very powerful when the underlying correlations between active and inactive covariates are unknown.

FarmSelect is applicable to a wide range of high dimensional sparse regression related problems that include but are not limited to linear model, generalized linear model, Gaussian graphic model, robust linear model and group LASSO. For the sparse linear regression, the proposed approach is equivalent to projecting the response variable and covariates onto the linear space orthogonal to the one spanned by the estimated factors. Existing algorithms that yield solution paths of LASSO can be directly applied in the second step. To demonstrate the finite sample performance of FarmSelect, we study two simulated and one empirical examples. The numerical results show FarmSelect can consistently select the true model even when the covariates are highly correlated while existing methods like LASSO, SCAD and elastic net fail to do so. An R-package FarmSelect (<https://cran.r-project.org/web/packages/FarmSelect>) is also provided to facilitate the implementation of our method.

Various methods have been studied to estimate the approximate factor model. Principal components analysis (PCA, [Stock and Watson, 2002](#)) is among one of the most popular ones. Data-driven estimation methods of the number of factors have been studied in extensive literature, such as [Bai and Ng \(2002\)](#), [Luo et al. \(2009\)](#), [Hallin and Liška \(2007\)](#), [Lam and Yao \(2012\)](#), and [Ahn and Horenstein \(2013\)](#) among others. Recently, a large amount of literature contributed to the asymptotic analysis of PCA under the ultra-high dimensional regime including [Johnstone and Lu \(2009\)](#), [Fan et al. \(2013\)](#), [Shen et al. \(2016\)](#) and [Wang and Fan \(2017\)](#), among others.

The rest of the paper is organized as follows. Section 2 overviews the problem setup including regularized M -estimator of sparse regression, the *irrepresentable condition* and approximate factor models. Section 3 introduces the model selection methodology of FarmSelect and studies the sparse generalized linear model as a showcase example. Some issues related to the estimation of approxi-

mate factor models will be discussed in Section 3 as well. Section 4 presents the general theoretical results. Section 5 provides simulation studies and Section 6 studies the forecast of U.S. bond risk premia. The appendix contains the technical proofs.

Here are some notations that will be used throughout the paper. \mathbf{I}_n denotes the $n \times n$ identity matrix; $\mathbf{0}$ refers to the $n \times m$ zero matrix; $\mathbf{0}_n$ and $\mathbf{1}_n$ represent the all-zero and all-one vectors in \mathbb{R}^n , respectively. For a matrix \mathbf{M} , we denote its matrix entry-wise max norm as $\|\mathbf{M}\|_{\max} = \max_{i,j} |M_{ij}|$ and denote by $\|\mathbf{M}\|_F$ and $\|\mathbf{M}\|_p$ its Frobenius and induced p -norms, respectively. $\lambda_{\min}(\mathbf{M})$ denotes the minimum eigenvalue of \mathbf{M} if it is symmetric. For $\mathbf{M} \in \mathbb{R}^{n \times m}$, $I \subseteq [n]$ and $J \subseteq [m]$, define $\mathbf{M}_{IJ} = (\mathbf{M}_{ij})_{i \in I, j \in J}$, $\mathbf{M}_I = (\mathbf{M}_{ij})_{i \in I, j \in [m]}$ and $\mathbf{M}_{\cdot J} = (\mathbf{M}_{ij})_{i \in [n], j \in J}$. For a vector $\mathbf{v} \in \mathbb{R}^p$ and $S \subseteq [p]$, define $\mathbf{v}_S = (\mathbf{v}_i)_{i \in S}$ to be its subvector. Let ∇ and ∇^2 be the gradient and Hessian operators. For $f : \mathbb{R}^p \rightarrow \mathbb{R}$ and $I, J \subseteq [p]$, define $\nabla_I f(x) = (\nabla f(x))_I$ and $\nabla_{IJ}^2 f(x) = (\nabla^2 f(x))_{IJ}$. $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ refers to the normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

2 Problem Setup

2.1 Regularized M -estimator

Let us begin with a family of high dimensional sparse regression problems in the following settings. From now on we suppose that $\{\mathbf{x}_t\}_{t=1}^n$ are $(p-1)$ -dimensional random vectors of covariates with zero mean¹, and $\{y_t\}_{t=1}^n$ are responses with each y_t sampled from some probability distribution $\mathbb{P}(z_t)$ parametrized by $z_t = \beta_0^* + \sum_{j=1}^{p-1} \beta_j^* \mathbf{x}_{tj} = (\mathbf{1}, \mathbf{x}_t^T) \boldsymbol{\beta}^*$. Here $\boldsymbol{\beta}^* = (\beta_0^*, \dots, \beta_{p-1}^*)^T \in \mathbb{R}^p$ is a sparse vector with $s \ll p$ non-zero elements. Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times (p-1)}$ and $\mathbf{y} = (y_1, \dots, y_n)^T \in \mathbb{R}^n$ be the design matrix and response vector, respectively. Define $\mathbf{X}_1 = (\mathbf{1}_n, \mathbf{X}) \in \mathbb{R}^{n \times p}$, where the subscript 1 refers to the all-one column added to the original design matrix \mathbf{X} .

Let $L_n(\mathbf{y}, \mathbf{X}_1 \boldsymbol{\beta})$ be some convex and differentiable loss function assigning a cost to any parameter $\boldsymbol{\beta} \in \mathbb{R}^p$. Suppose that $\boldsymbol{\beta}^*$ is the unique minimizer of the population risk $\mathbb{E}[L_n(\mathbf{y}, \mathbf{X}_1 \boldsymbol{\beta})]$. Under the high-dimensional regime, it is natural to estimate $\boldsymbol{\beta}^*$ via a regularized M -estimator as follows:

$$\tilde{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\operatorname{argmin}} \{L_n(\mathbf{y}, \mathbf{X}_1 \boldsymbol{\beta}) + \lambda R_n(\boldsymbol{\beta})\}, \quad (2.1)$$

where $R_n : \mathbb{R}^p \rightarrow \mathbb{R}_+$ is a norm that penalizes the use of a nonsparse vector $\boldsymbol{\beta}$ and $\lambda > 0$ is a tuning

¹We use $(p-1)$ instead of p to denote the number of covariates so that there are p coefficients including the intercept. In addition, we center the covariates if they could have non-zero means. Whether this step is done or not does not affect the estimation of $\{\boldsymbol{\beta}_j^*\}_{j=1}^p$, but does affect the intercept β_0^* .

parameter.

A special case of this problem is the L_1 penalized likelihood estimation of generalized linear models. Suppose the conditional density function of Y given covariates \mathbf{x} is a member of the exponential family, i.e.

$$f(y|\mathbf{x}, \boldsymbol{\beta}^*) \propto \exp[yz - b(z) + c(y)], \quad (2.2)$$

where $z = \beta_0^* + \sum_{j=1}^{p-1} \beta_j^* x_j = (1, \mathbf{x}^T) \boldsymbol{\beta}^*$, $b(\cdot)$ and $c(\cdot)$ are known functions, and $\boldsymbol{\beta}^*$ is an unknown coefficient vector of interest. It is commonly assumed that $b(\cdot)$ is strictly convex. Taking the loss function to be the negative log-likelihood function and the penalty function to be the L_1 norm, the regularized M -estimator of $\boldsymbol{\beta}^*$ admits the form

$$\tilde{\boldsymbol{\beta}} \in \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{t=1}^n [-y_t(1, \mathbf{x}_t^T) \boldsymbol{\beta} + b((1, \mathbf{x}_t^T) \boldsymbol{\beta})] + \lambda \|\boldsymbol{\beta}\|_1 \right\}. \quad (2.3)$$

2.2 Irrepresentable condition

We expect a good estimator of (2.1) to achieve estimation as well as selection consistency. The former requires $\|\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\| \xrightarrow{P} \mathbf{0}$ for some norm $\|\cdot\|$ as $n \rightarrow \infty$; while the latter requires $P(\operatorname{supp}(\tilde{\boldsymbol{\beta}}) = \operatorname{supp}(\boldsymbol{\beta}^*)) \rightarrow 1$ as $n \rightarrow \infty$. In general, the estimation consistency does not imply the selection consistency and vice versa. To study the selection consistency, we consider a stronger condition named general sign consistency as follows.

Definition 2.1 (Sign consistency). *An estimate $\tilde{\boldsymbol{\beta}}$ is sign consistent with respect to $\boldsymbol{\beta}^*$ if $\lim_{n \rightarrow \infty} P(\operatorname{sign}(\tilde{\boldsymbol{\beta}}) = \operatorname{sign}(\boldsymbol{\beta}^*)) = 1$.*

Zhao and Yu (2006) studied the LASSO estimator and showed there exists an *irrepresentable condition* which is sufficient and almost necessary for both sign and estimation consistencies for sparse linear model. Without loss of generality, we assume $\operatorname{supp}(\boldsymbol{\beta}^*) = [s] = S$. Denote $(\mathbf{X}_1)_S$ and $(\mathbf{X}_1)_{S^c}$ as the submatrices of \mathbf{X}_1 defined by its first s columns and the rest $(p - s)$ columns, respectively. Then the *irrepresentable condition* requires some $\tau \in (0, 1)$, such that

$$\|(\mathbf{X}_1)_{S^c}^T (\mathbf{X}_1)_S [(\mathbf{X}_1)_S^T (\mathbf{X}_1)_S]^{-1}\|_\infty \leq 1 - \tau. \quad (2.4)$$

For general regularized M -estimator (2.1) to achieve both sign and estimation consistencies, Lee et al. (2015) proposed a generalized *irrepresentable condition*. When applied to the L_1 regularizer, it becomes

$$\|\nabla_{S^c S}^2 L(\boldsymbol{\beta}^*) [\nabla_{SS}^2 L(\boldsymbol{\beta}^*)]^{-1}\|_\infty \leq 1 - \tau, \quad (2.5)$$

for some $\tau \in (0, 1)$, where $L(\boldsymbol{\beta}) = L_n(\mathbf{y}, \mathbf{X}_1\boldsymbol{\beta})$. It is easy to check (2.5) is equivalent to (2.4) under the LASSO case. The generalized *irrepresentable condition* will easily get violated when there exists strong correlations between active and inactive variables. Even if it holds, the key parameter τ can be very close to zero, making it hard to select the correct model and obtain small estimation errors simultaneously.

2.3 Approximate factor model

To go beyond the assumption on weakly correlation, a natural extension is conditional weak correlation. Suppose covariates are dependent through latent common factors. Given these common factors, the idiosyncratic components are weakly correlated. Factor model has been well studied in econometrics and statistics literature, we refer to Lawley and Maxwell (1971); Stock and Watson (2002); Bai and Ng (2002); Forni et al. (2013); Fan et al. (2013), among others. For an overview, see Fan et al. (2019+).

We assume that $\{\mathbf{x}_t\}_{t=1}^n \subseteq \mathbb{R}^{p-1}$ follows the approximate factor model

$$\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t, \quad t \in [n], \quad (2.6)$$

where $\{\mathbf{f}_t\}_{t=1}^n \subseteq \mathbb{R}^K$ are latent factors, $\mathbf{B} \in \mathbb{R}^{(p-1) \times K}$ is a loading matrix, and $\{\mathbf{u}_t\}_{t=1}^n \subseteq \mathbb{R}^{p-1}$ are idiosyncratic components. Note that \mathbf{x}_t is the only observable quantity. Throughout the paper, K is assumed to be independent of n , which is frequently imposed in the literature of factor model (Fan et al., 2013). We assume that $\{\mathbf{f}_t, \mathbf{u}_t\}_{t=1}^n$ come from a time series $\{\mathbf{f}_t, \mathbf{u}_t\}_{t=-\infty}^{\infty}$. Denote $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_n)^T \in \mathbb{R}^{n \times K}$ and $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_n)^T \in \mathbb{R}^{n \times (p-1)}$. Then (2.6) can be written in a more compact form:

$$\mathbf{X} = \mathbf{F}\mathbf{B}^T + \mathbf{U}. \quad (2.7)$$

We impose the following identifiability assumption (Fan et al., 2013). Here we only put the most basic assumption for factor model, and more can be found in Section 3.3 where estimation of factor model is discussed.

Assumption 2.1. Assume that $\text{cov}(\mathbf{f}_t) = \mathbf{I}_K$, $\mathbf{B}^T\mathbf{B}$ is diagonal, and all the eigenvalues of $\mathbf{B}^T\mathbf{B}/p$ are bounded away from 0 and ∞ as $p \rightarrow \infty$.

3 Factor-adjusted regularized model selection

3.1 Methodology

To illustrate the main idea, we temporarily assume \mathbf{f}_t and \mathbf{u}_t to be observable. Define $\mathbf{B}_0 = (\mathbf{0}_K, \mathbf{B}^T)^T \in \mathbb{R}^{K \times p}$ and $\mathbf{U}_1 = (\mathbf{1}_n, \mathbf{U}) \in \mathbb{R}^{n \times p}$. By the approximate factor model (2.7), we have decompositions $\mathbf{X}_1 = \mathbf{F}\mathbf{B}_0^T + \mathbf{U}_1$ and

$$\mathbf{X}_1\boldsymbol{\beta} = \mathbf{F}\mathbf{B}_0^T\boldsymbol{\beta} + \mathbf{U}_1\boldsymbol{\beta} = \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}_1\boldsymbol{\beta},$$

where $\boldsymbol{\gamma} = \mathbf{B}_0^T\boldsymbol{\beta} \in \mathbb{R}^K$. The regularized M -estimator (2.1) can be written as

$$\tilde{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} = \mathbf{B}_0^T\boldsymbol{\beta} \in \mathbb{R}^K}{\operatorname{argmin}} \{L_n(\mathbf{y}, \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}_1\boldsymbol{\beta}) + \lambda R_n(\boldsymbol{\beta})\}.$$

Instead of using $\tilde{\boldsymbol{\beta}}$ to estimate $\boldsymbol{\beta}^*$, we regard $\boldsymbol{\gamma}$ as nuisance parameters, drop the constraint $\boldsymbol{\gamma} = \mathbf{B}_0^T\boldsymbol{\beta}$, and consider a new estimator

$$\hat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}^K}{\operatorname{argmin}} \{L_n(\mathbf{y}, \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}_1\boldsymbol{\beta}) + \lambda R_n(\boldsymbol{\beta})\}, \quad (3.1)$$

namely $(\mathbf{u}_t^T, \mathbf{f}_t^T)^T$ are now regarded as new covariates. In other words, by lifting the covariate space from \mathbb{R}^p to \mathbb{R}^{p+K} , the highly dependent covariates \mathbf{x}_t are replaced by weakly dependent ones.

The theory for us to ignore the constraint $\boldsymbol{\gamma} = \mathbf{B}_0^T\boldsymbol{\beta}$ is given by the following lemma, whose proof is given by Appendix A.

Lemma 3.1. *Consider the generalized linear model (2.2), let $L_n(\mathbf{y}, \mathbf{z}) = \frac{1}{n} \sum_{t=1}^n [-y_t z_t + b(z_t)]$, $\eta_t = y_t - b'((1, \mathbf{x}_t^T)\boldsymbol{\beta}^*)$ and $\mathbf{w}_t = (1, \mathbf{u}_t^T, \mathbf{f}_t^T)^T$. If $E(\eta_t \mathbf{w}_t) = \mathbf{0}_{p+K}$, then*

$$(\boldsymbol{\beta}^*, \mathbf{B}_0^T\boldsymbol{\beta}^*) = \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}^K}{\operatorname{argmin}} E[L_n(\mathbf{y}, \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}_1\boldsymbol{\beta})].$$

It is worth pointing out that the assumption $E(\eta_t \mathbf{w}_t) = \mathbf{0}_{p+K}$ is very mild and natural. We just assume the residual η_t and augmented covariates \mathbf{w}_t to be uncorrelated, which is almost as weak as the standard condition $E(\eta_t | \mathbf{x}_t) = 0$ for the generalized linear model. For example, in the linear model $y_t = (1, \mathbf{x}_t^T)\boldsymbol{\beta}^* + \eta_t$, we strengthen the condition only from $E(\eta_t \mathbf{x}_t) = 0$ to $E(\eta_t \mathbf{f}_t) = 0$ and $E(\eta_t \mathbf{u}_t) = 0$. In particular, the assumptions hold if η_t is independent of \mathbf{u}_t and \mathbf{f}_t .

By construction, (\mathbf{U}, \mathbf{F}) has much weaker cross-sectional correlation than \mathbf{X} . Thus, the penalized profile likelihood (3.1) removes the effect of strong correlations caused by the latent factors. It can be implemented as follows:

Step 1: Initial estimation. Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ be the design matrix. Fit the approximate factor model (2.7) and denote $\widehat{\mathbf{B}}$, $\widehat{\mathbf{F}}$ and $\widehat{\mathbf{U}} = \mathbf{X} - \widehat{\mathbf{F}}\widehat{\mathbf{B}}^T$ the obtained estimators of \mathbf{B} , \mathbf{F} and \mathbf{U} respectively by using the principal component analysis (Bai, 2003; Fan et al., 2013; Fan et al., 2019+). More specifically, the columns of $\widehat{\mathbf{F}}/\sqrt{n}$ are the eigenvectors of $\mathbf{X}\mathbf{X}^T$ corresponding to the top K eigenvalues, $\widehat{\mathbf{B}} = n^{-1}\mathbf{X}^T\widehat{\mathbf{F}}$. This is the same as $\widehat{\mathbf{B}} = (\sqrt{\lambda_1}\boldsymbol{\xi}_1, \dots, \sqrt{\lambda_K}\boldsymbol{\xi}_K)$ and $\widehat{\mathbf{F}} = \mathbf{X}\widehat{\mathbf{B}}\text{diag}(\lambda_1^{-1} \dots, \lambda_K^{-1})$, where $\{\lambda_j\}_{j=1}^K$ and $\{\boldsymbol{\xi}_j\}_{j=1}^K$ are top K eigenvalues in descending order and their associated eigenvectors of the sample covariance matrix.

Step 2: Augmented M-estimation. Define $\widehat{\mathbf{W}} = (\mathbf{1}_n, \widehat{\mathbf{U}}, \widehat{\mathbf{F}}) \in \mathbb{R}^{n \times (p+K)}$ and $\boldsymbol{\theta} = (\boldsymbol{\beta}^T, \boldsymbol{\gamma}^T)^T \in \mathbb{R}^{p+K}$. Then $\widehat{\boldsymbol{\beta}}$ is obtained from the first p entries of the solution to the augmented problem

$$\widehat{\boldsymbol{\theta}} \in \underset{\boldsymbol{\theta} \in \mathbb{R}^{p+K}}{\text{argmin}} \left\{ L_n(\mathbf{y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda R_n(\boldsymbol{\theta}_{[p]}) \right\}. \quad (3.2)$$

We call the above two-step method as the factor-adjust regularized model selection (*FarmSelect*). If \mathbf{u}_t is independent of \mathbf{f}_t and the variables in the idiosyncratic component \mathbf{u}_t are weakly correlated, then the columns in $\widehat{\mathbf{W}} = (\mathbf{1}_n, \widehat{\mathbf{U}}, \widehat{\mathbf{F}})$ are weakly correlated as long as \mathbf{F} and \mathbf{U} are well estimated. Hence, we successfully transform the problem from model selection with highly correlated covariates \mathbf{X} in (2.1) to model selection with weakly correlated or uncorrelated ones by lifting the space to higher dimension. The augmented problem (3.2) is a convex optimization problem which can be minimized via many existing convex optimization algorithms, such as coordinate descent (e.g. Friedman et al., 2010) and ADMM (Boyd et al., 2011).

3.2 Example: sparse linear model

Now we illustrate the FarmSelect procedure using sparse linear regression, where $\mathbf{y} = \mathbf{X}_1\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}$. With aforementioned notation, we have

$$\mathbf{y} = \mathbf{X}_1\boldsymbol{\beta}^* + \boldsymbol{\varepsilon} = \widehat{\mathbf{F}}\widehat{\mathbf{B}}_0^T\boldsymbol{\beta}^* + \widehat{\mathbf{U}}_1\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}. \quad (3.3)$$

The augmented M -estimator (3.2) for the sparse linear model is of the following form:

$$\widehat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}^K}{\text{argmin}} \left\{ \frac{1}{2n} \|\mathbf{y} - \widehat{\mathbf{F}}\boldsymbol{\gamma} - \widehat{\mathbf{U}}_1\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1 \right\}.$$

Solving the least-squares problem with respect to $\boldsymbol{\gamma}$, we have the penalized profile least-squares solution

$$\widehat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\text{argmin}} \left\{ \frac{1}{2n} \|(\mathbf{I}_n - \widehat{\mathbf{P}})(\mathbf{y} - \widehat{\mathbf{U}}_1\boldsymbol{\beta})\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1 \right\}, \quad (3.4)$$

where $\widehat{\mathbf{P}} = \widehat{\mathbf{F}}(\widehat{\mathbf{F}}^T \widehat{\mathbf{F}})^{-1} \widehat{\mathbf{F}}^T$ is the $n \times n$ projection matrix onto the column space of $\widehat{\mathbf{F}}$. As the decorrelation step does not depend on the choice of the regularizer $R(\cdot)$, FarmSelect can be applied to a wide range of penalized least squares problems such as SCAD, group LASSO, elastic net, fused LASSO, other folded concave penalties, and so on.

There is another way to understand this method. By left multiplying the projection matrix $(\mathbf{I}_n - \widehat{\mathbf{P}})$ to both sides of (3.3), we have

$$(\mathbf{I}_n - \widehat{\mathbf{P}})\mathbf{y} = (\mathbf{I}_n - \widehat{\mathbf{P}})\widehat{\mathbf{U}}_1\boldsymbol{\beta}^* + (\mathbf{I}_n - \widehat{\mathbf{P}})\boldsymbol{\varepsilon}, \quad (3.5)$$

where $(\mathbf{I}_n - \widehat{\mathbf{P}})\widehat{\mathbf{U}}_1$ can be treated as the decorrelated design matrix and $(\mathbf{I}_n - \widehat{\mathbf{P}})\mathbf{y}$ is the corresponding response variable. From (3.5) we see that the method in Kneip and Sarda (2011) coincides with FarmSelect in the linear case. However, the projection-based representation only makes sense in sparse linear regression. In contrast, our idea of profile likelihood directly generalizes to more general problems.

3.3 Estimating factor models

Principal component analysis (PCA) is frequently used to estimate latent factors for model (2.7). The estimated matrix of latent factors $\widehat{\mathbf{F}}$ is \sqrt{n} times the eigenvectors corresponding to the K largest eigenvalues of the $n \times n$ matrix $\mathbf{X}\mathbf{X}^T$. Using the normalization $\mathbf{F}^T \mathbf{F}/n = \mathbf{I}_K$ yields $\widehat{\mathbf{B}} = \mathbf{X}^T \widehat{\mathbf{F}}/n$. Now we introduce the asymptotic properties of estimated factors and idiosyncratic components. We adopt the regularity assumptions in Fan et al. (2013), which are similar to the ones in Bai (2003) and other literature on high-dimensional factor analysis.

Assumption 3.1. 1. $\{\mathbf{f}_t, \mathbf{u}_t\}_{t=1}^\infty$ is strictly stationary. In addition, $\mathbb{E} f_{tk} = \mathbb{E} u_{tj} = \mathbb{E}(u_{tj} f_{tk}) = 0$ for all $i \in [n]$, $j \in [p-1]$ and $k \in [K]$;

2. There exist constants $c_1, c_2 > 0$ such that $\lambda_{\min}(\text{cov}(\mathbf{u}_t)) > c_1$, $\|\text{cov}(\mathbf{u}_t)\|_1 < c_2$ and $\min_{j,k \in [p-1]} \text{var}(u_{tj} u_{tk}) > c_1$;

3. There exist $r_1, r_2 > 0$ and $b_1, b_2 > 0$ such that for any $s > 0$, $j \in [p-1]$ and $k \in [K]$, $\mathbb{P}(|u_{tj}| > s) \leq \exp(-(s/b_1)^{r_1})$ and $\mathbb{P}(|f_{tk}| > s) \leq \exp(-(s/b_2)^{r_2})$.

Assumption 3.2. Let $\mathcal{F}_{-\infty}^0$ and \mathcal{F}_T^∞ denote the σ -algebras generated by $\{(\mathbf{f}_t, \mathbf{u}_t) : i \leq 0\}$ and $\{(\mathbf{f}_t, \mathbf{u}_t) : i \geq T\}$ respectively. Assume the existence of $r_3, C > 0$ such that $3/r_1 + 3/(2r_2) + 1/r_3 > 1$

and for all $T \geq 1$,

$$\sup_{A \in \mathcal{F}_{-\infty}^0, B \in \mathcal{F}_T^\infty} |\mathbf{P}(A)\mathbf{P}(B) - \mathbf{P}(AB)| \leq \exp(-CT^{r_3});$$

Assumption 3.3. *There exists $M > 0$ such that for all $t, s \in [n]$, we have $\|\mathbf{B}\|_{\max} < M$, $\mathbf{E}\{p^{-1/2}[\mathbf{u}_t^T \mathbf{u}_s - \mathbf{E}(\mathbf{u}_t^T \mathbf{u}_s)]^4\} < M$ and $\mathbf{E} \|p^{-1/2} \mathbf{B}^T \mathbf{u}_t\|_2^4 < M$.*

We summarize useful properties of $\hat{\mathbf{F}}$ and $\hat{\mathbf{U}}$ in Lemma 3.2, which directly follows from Lemmas 10-12 in Fan et al. (2013).

Lemma 3.2. *Let $\gamma^{-1} = 3/r_1 + 3/(2r_2) + 1/r_3 + 1$. Suppose that $\log p = o(n^{\gamma/6})$, $n = o(p^2)$, and Assumptions 2.1, 3.1, 3.2 and 3.3 hold. There exists a nonsingular matrix $\mathbf{H}_0 \in \mathbb{R}^{K \times K}$ such that*

1. $\|\hat{\mathbf{F}}\mathbf{H}_0 - \mathbf{F}\|_{\max} = O_{\mathbf{P}}(\frac{1}{\sqrt{n}} + \frac{n^{1/4}}{\sqrt{p}});$
2. $\max_{k \in [K]} n^{-1} \sum_{t=1}^n |(\hat{\mathbf{F}}\mathbf{H}_0)_{jk} - f_{tk}|^2 = O_{\mathbf{P}}(\frac{1}{n} + \frac{1}{p})$
3. $\max_{j \in [p-1]} n^{-1} \sum_{t=1}^n |\hat{u}_{ji} - u_{ji}|^2 = O_{\mathbf{P}}(\frac{\log p}{n} + \frac{1}{p});$
4. $\|\hat{\mathbf{U}} - \mathbf{U}\|_{\max} = o_{\mathbf{P}}(1).$

A practical issue arises on how to choose the number of factors, i.e. K . As latent factors, loading and idiosyncratic components are all unobservable in the approximate factor model, the estimation of K is an intrinsic unsupervised learning problem. From the inference point of view, existing literature (Chamberlain and Rothschild, 1982; Stock and Watson, 2002; Bai and Ng, 2002, among others) usually assumes that there exists a non-negative integer K such that the first K population eigenvalues of \mathbf{X} are diverging with p , while the rest $p - K$ eigenvalues are bounded. From the dimension reduction point of view, selecting K is to find a proper trade-off between goodness-of-fit and compactness of model. In this paper, we follow a conditional sparsity perspective (Fan et al., 2013) regarding the role of K , where K is the smallest non-negative integer such that the idiosyncratic components $\mathbf{U} = \mathbf{X} - \mathbf{F}\mathbf{B}^T$ is weakly correlated. In this regard, for our purpose of model selection, small overestimation of K does not seriously affect the adjusted model selection.

We adopt the modified ratio method, e.g. equation (10) in Chang et al. (2015), for the numerical studies in this paper due to its simplicity. Let $\lambda_k(\mathbf{X}\mathbf{X}^T)$ be the k th largest eigenvalue of $\mathbf{X}\mathbf{X}^T$, K_{\max} be a prescribed upper bound and C_n be a constant that depends on n and p . The number of factors can be estimated by

$$\hat{K} = \operatorname{argmin}_{k \leq K_{\max}} \frac{\lambda_{k+1}(\mathbf{X}\mathbf{X}^T) + C_n}{\lambda_k(\mathbf{X}\mathbf{X}^T) + C_n} \quad (3.6)$$

for some given C_n . When \mathbf{X} itself is weakly correlated, one can estimate K as 0.

Beside the modified ratio method, [Bai and Ng \(2002\)](#) studied the convergence and consistency estimation of K for high dimensional factor models. They proposed to estimate K by minimizing a family of information criteria. We refer to equation (9) in [Bai and Ng \(2002\)](#) for viable examples.

3.4 Factor-adjusted variable screening

Screening methods (e.g. [Fan and Lv, 2008](#); [Fan and Song, 2009](#); [Wang and Leng, 2016](#)) are computationally attractive and thus popular for ultra-high dimensional data analysis. However, the screening methods tend to include too many variables when there exist strong correlations among covariates ([Fan and Lv, 2008](#); [Wang and Leng, 2016](#)). As an extension of FarmSelect, we propose the following conditional variable screening method to tackle this problem.

Step 1: Initial estimation. We fit the approximate factor model (2.7) to obtain $\hat{\mathbf{B}}$, $\hat{\mathbf{F}}$ and $\hat{\mathbf{U}}$.

Step 2: Augmented marginal regression. For $j \in [p - 1]$, let $\hat{\mathbf{U}}_j$ be the j -th column of $\hat{\mathbf{U}}$ and compute

$$(\hat{\alpha}_j, \hat{\beta}_j, \hat{\gamma}_j) = \underset{\alpha \in \mathbb{R}, \beta \in \mathbb{R}, \gamma \in \mathbb{R}^K}{\operatorname{argmin}} L_n(\mathbf{y}, \mathbf{1}_n \alpha + \hat{\mathbf{U}}_j \beta + \hat{\mathbf{F}} \gamma). \quad (3.7)$$

Step 3 Screening. Return $\{j : |\hat{\beta}_j| \geq \xi\}$ for some prescribed threshold ξ .

For sparse linear regression, our screening method reduces to the factor-profiled screening method proposed by [Wang \(2012\)](#).

4 Theoretical results

4.1 FarmSelect with approximate factor model

Now we establish theoretical guarantees of the FarmSelect estimator (3.2). Recall that β^* is equal to the first p entries of θ^* . Define $S = \operatorname{supp}(\theta^*)$, $S_1 = \operatorname{supp}(\beta^*)$ and $S_2 = [p + K] \setminus S$. When the covariates \mathbf{X} admit the approximate factor model (2.7), the oracle procedure uses true augmented covariates $\mathbf{w}_t = (1, \mathbf{u}_t^T, \mathbf{f}_t^T)^T$ for $t \in [n]$ and solves

$$\min_{\theta} \{L_n(\mathbf{y}, \mathbf{W}\theta) + \lambda \|\theta_{[p]}\|_1\},$$

where $\mathbf{W} = (\mathbf{w}_1^T, \dots, \mathbf{w}_n^T)^T = (\mathbf{U}_1, \mathbf{F})$. However, \mathbf{W} is not observable in practice. Hence we need to use its estimator $\widehat{\mathbf{W}}$ and solve

$$\min_{\boldsymbol{\theta}} \{L_n(\mathbf{y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[p]}\|_1\}.$$

Below the error induced by the factor estimation will be studied carefully. To deliver a clear discussion on the conditions and results, we focus on the FarmSelect estimator for the generalized linear model (2.3), and assume that the covariates are generated from the approximate factor model (2.7).

Assumption 4.1 (Smoothness). $b(z) \in C^3(\mathbb{R})$. For some constants M_2 and M_3 , we have $0 \leq b''(z) \leq M_2$ and $|b'''(z)| \leq M_3$, $\forall z$.

Assumption 4.2 (Restricted strong convexity and irrepresentable condition). Let $\boldsymbol{\theta}^* = \begin{pmatrix} \boldsymbol{\beta}^* \\ \mathbf{B}_0^T \boldsymbol{\beta}^* \end{pmatrix}$.

Assume the existence of $\kappa_2 > \kappa_\infty > 0$ and $\tau \in (0, 1)$ such that

$$\begin{aligned} \|\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)^{-1}\|_\ell &\leq \frac{1}{4\kappa_\ell}, \quad \text{for } \ell = 2 \text{ and } \infty, \\ \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)[\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)^{-1}]\|_\infty &\leq 1 - 2\tau. \end{aligned} \tag{4.1}$$

Assumption 4.3 (Estimation of factor model). $\|\mathbf{W}\|_{\max} \leq \frac{M_0}{2}$ for some constant $M_0 > 0$. In addition, there exist $K \times K$ nonsingular matrix \mathbf{H}_0 , and $\mathbf{H} = \begin{pmatrix} \mathbf{I}_p & \mathbf{0}_{p \times K} \\ \mathbf{0}_{K \times p} & \mathbf{H}_0 \end{pmatrix}$ such that for $\overline{\mathbf{W}} = \widehat{\mathbf{W}}\mathbf{H}$, we have $\|\overline{\mathbf{W}} - \mathbf{W}\|_{\max} \leq \frac{M_0}{2}$ and $\max_{j \in [p+K]} \left(\frac{1}{n} \sum_{t=1}^n |\overline{w}_{tj} - w_{tj}|^2 \right)^{1/2} \leq \frac{2\kappa_\infty \tau}{3M_0 M_2 |S|}$.

Before presenting the main results, we make a few remarks on the assumptions.

1. Assumption 4.1 holds for a large family of generalized linear models. For example, linear model has $b(z) = \frac{1}{2}z^2$, $M_2 = 1$ and $M_3 = 0$; logistic model has $b(z) = \log(1 + e^z)$ and finite M_2, M_3 .
2. The first inequality in (4.1) involves only a small matrix and holds easily, and the second inequality there is related to the generalized irrepresentable condition. Standard concentration inequalities (e.g. the Bernstein inequality for weakly dependent variables in Merlevède et al. (2011)) yield that Assumption 4.2 holds with high probability as long as $\mathbb{E}[\nabla^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)]$ satisfies similar conditions.

3. Here we show an example where the irrerepresentable condition holds for the augmented covariates \mathbf{W} but fails to hold for the original ones \mathbf{X} . Suppose the covariates $\{\mathbf{x}_i\}_{i=1}^n$ are generated from a single factor model $\mathbf{x}_i = \mathbf{b}f_i + \mathbf{u}_i$, where $\mathbf{b} = (1, 1, 2, \dots, 2)^T \in \mathbb{R}^p$, $f_i \sim N(0, 1)$, $\mathbf{u}_i \sim N(\mathbf{0}, \mathbf{I}_p)$ and it is independent of f_i . Consider a sparse linear model $y_i = \mathbf{x}_i^T \boldsymbol{\beta}^* + \varepsilon_i$, where $\{\varepsilon_i\}_{i=1}^n$ are i.i.d. $N(0, \sigma^2)$ variables for some $\sigma > 0$ and they are independent of $\{\mathbf{f}_i, \mathbf{u}_i\}_{i=1}^n$; $\boldsymbol{\beta}^* = (1, 1, 0, \dots, 0)^T \in \mathbb{R}^p$. Let $L_n(\mathbf{y}, \mathbf{z}) = \frac{1}{2n} \|\mathbf{y} - \mathbf{z}\|_2^2$. Note that $\mathbf{w}_i = (\mathbf{u}_i^T, f_i)^T \in \mathbb{R}^{p+1}$, $\boldsymbol{\theta}^* = ((\boldsymbol{\beta}^*)^T, \mathbf{b}^T \boldsymbol{\beta}^*)^T \in \mathbb{R}^{p+1}$,

$$y_i = \mathbf{x}_i^T \boldsymbol{\beta}^* + \varepsilon_i = f_i \mathbf{b}^T \boldsymbol{\beta}^* + \mathbf{u}_i^T \boldsymbol{\beta}^* + \varepsilon_i = \mathbf{w}_i^T \boldsymbol{\theta}^* + \varepsilon_i,$$

and $\nabla^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}) = \frac{1}{n} \mathbf{W}^T \mathbf{W}$. Thanks to the independence, we have $\mathbf{w}_i \sim N(\mathbf{0}, \mathbf{I}_{p+1})$ and $\mathbb{E} \nabla^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}) = \mathbf{I}_{p+1}$. Hence

$$\nabla_{S_2 S}^2 \mathbb{E} L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*) [\nabla_{S S}^2 \mathbb{E} L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)]^{-1} = \mathbf{0},$$

where $S = \{1, 2, p+1\}$ and $S_2 = \{3, \dots, p\}$. With high probability, the irrerepresentable condition also holds for the empirical quantity $L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)$. On the other hand, we observe that $\mathbf{x}_i \sim N(\mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = \mathbf{b}\mathbf{b}^T + \mathbf{I}_p$. When \mathbf{X} are used as covariates, we have $S = \{1, 2\}$ and $S_2 = \{3, \dots, p\}$. From $\boldsymbol{\Sigma}_{SS} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ and $\boldsymbol{\Sigma}_{S^c S} = 2 \cdot \mathbf{1}_{(p-2) \times 2}$, we get $\boldsymbol{\Sigma}_{S^c S} \boldsymbol{\Sigma}_{SS}^{-1} = \frac{2}{3} \cdot \mathbf{1}_{(p-2) \times 2}$ and $\|\boldsymbol{\Sigma}_{S^c S} \boldsymbol{\Sigma}_{SS}^{-1}\|_\infty = \frac{4}{3} > 1$. Hence the irrerepresentable condition is violated.

4. Under the conditions of Lemma 3.2, we have $\|\bar{\mathbf{W}} - \mathbf{W}\|_{\max} = o_P(1)$ and $\max_{j \in [p+K]} \left(\frac{1}{n} \sum_{t=1}^n |\bar{w}_{tj} - w_{tj}|^2 \right)^{1/2} = O_P\left(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}}\right)$, where $\bar{\mathbf{W}} = \widehat{\mathbf{W}}\mathbf{H}$, $\mathbf{H} = \begin{pmatrix} \mathbf{I}_p & \mathbf{0}_{p \times K} \\ \mathbf{0}_{K \times p} & \mathbf{H}_0 \end{pmatrix}$ and some proper \mathbf{H}_0 . Hence $|S|^2 \left(\frac{\log p}{n} + \frac{1}{p}\right) = O(1)$ can guarantee Assumption 4.3 to hold with high probability.

Theorem 4.1. Suppose Assumptions 4.1-4.3 hold. Define $M = M_0^3 M_3 |S|^{3/2}$ and

$$\varepsilon = \max_{j \in [p+K]} \left| \frac{1}{n} \sum_{t=1}^n \bar{w}_{tj} [-y_t + b'((1, \mathbf{x}_t^T) \boldsymbol{\beta}^*)] \right|.$$

If $\frac{7\varepsilon}{\tau} < \lambda < \frac{\kappa_2 \kappa_\infty \tau}{12M\sqrt{|S|}}$, then we have $\text{supp}(\widehat{\boldsymbol{\beta}}) \subseteq \text{supp}(\boldsymbol{\beta}^*)$ and

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_\infty \leq \frac{6\lambda}{5\kappa_\infty}, \quad \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 \leq \frac{4\lambda\sqrt{|S|}}{\kappa_2}, \quad \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_1 \leq \frac{6\lambda|S|}{5\kappa_\infty}.$$

In addition, if $\varepsilon < \frac{\kappa_2 \kappa_\infty \tau^2}{12CM\sqrt{|S|}}$ and $\min\{|\boldsymbol{\beta}_j^*| : \boldsymbol{\beta}_j^* \neq 0, j \in [p]\} > \frac{6C\varepsilon}{5\kappa_\infty \tau}$ hold for some $C > 7$, then by taking $\lambda \in (\frac{7\varepsilon}{\tau}, \frac{C\varepsilon}{\tau})$ we can achieve the sign consistency $\text{sign}(\widehat{\boldsymbol{\beta}}) = \text{sign}(\boldsymbol{\beta}^*)$.

By taking $\lambda \asymp \varepsilon$, one can achieve the sign consistency and $\|\widehat{\beta} - \beta^*\|_\infty / \varepsilon = O_P(1)$, $\|\widehat{\beta} - \beta^*\|_2 / \varepsilon = O_P(\sqrt{|S|})$ and $\|\widehat{\beta} - \beta^*\|_1 / \varepsilon = O_P(|S|)$. Hence ε is a key quantity characterizing the error rate of our FarmSelect estimator, whose size is controlled using the following lemma.

Lemma 4.1. *Let $\eta_t = y_t - b'((1, \mathbf{x}_t^T)\beta^*)$ and $\mathbf{w}_t = (1, \mathbf{u}_t^T, \mathbf{f}_t^T)^T$. Assume that $\{\mathbf{w}_t, \eta_t\}_{-\infty}^\infty$ is strictly stationary and satisfies the following conditions*

1. $E(\eta_t \mathbf{w}_t) = \mathbf{0}$;
2. *There exist constants $b, \gamma_1 > 0$ such that $P(|\eta_t| > s) \leq \exp(1 - (s/b)^{\gamma_1})$ for all $t \in \mathbb{Z}$ and $s \geq 0$;*
3. *There exist constants $c, \gamma_3 > 0$ such that for all $T \geq 1$,*

$$\sup_{A \in \overline{\mathcal{F}}_{-\infty}^0, B \in \overline{\mathcal{F}}_T^\infty} |P(A)P(B) - P(AB)| \leq \exp(-cT^{\gamma_3}),$$

where $\overline{\mathcal{F}}_{-\infty}^0$ and $\overline{\mathcal{F}}_{-\infty}^0$ denote the σ -algebras generated by $\{(\mathbf{w}_t, \eta_t) : i \leq 0\}$ and $\{(\mathbf{w}_t, \eta_t) : i \geq T\}$ respectively;

In addition, suppose that the assumptions in Lemma 3.2 hold. Then we have

$$\varepsilon = \max_{j \in [p+K]} \left| \frac{1}{n} \sum_{t=1}^n \bar{w}_{tj} \eta_t \right| = O_P\left(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}}\right).$$

Recall that the assumption $E(\eta_t \mathbf{w}_t) = \mathbf{0}$ has been used in Lemma 3.1 as a cornerstone of our FarmSelect methodology. The rest in the list are standard conditions similar to Assumptions 3.1-3.3. All of them are mild and interpretable.

Lemma 4.1 asserts that $\varepsilon = O_P(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}})$. The first term $\sqrt{\frac{\log p}{n}}$ corresponds to the optimal rate of convergence for high-dimensional M -estimator (e.g. Bickel et al., 2009). The second term $\frac{1}{\sqrt{p}}$ is the price we pay for factor estimation, which is negligible if $n = O(p \log p)$. In that high-dimensional regime, all the error bounds for $\|\widehat{\beta} - \beta^*\|_\ell$ ($\ell = 1, 2, \infty$) match the optimal ones in the literature.

4.2 Factor-adjusted variable screening

In this subsection, we study the factor-adjusted variable screening procedure described in Section 3.4. The lemma below considers the population version of the factor-adjusted screening procedure.

It shows that as long as $E(u_j y) \neq 0$ for all j 's in the active set $\text{supp}(\beta^*)$, where u_j is the j -th idiosyncratic component, the screening retains all the important variables. Furthermore, if $E(u_j y) = 0$ for all j 's outside the active set, the screening procedure exactly recovers the active set.

Lemma 4.2. *Let \mathbf{f} be a K -dimensional random vector, u be a zero-mean random variable and is independent of \mathbf{f} , y be another random variable living in the same probability space, and $b \in C^2(\mathbb{R})$ such that $0 < b'' \leq M$. Assume that u , y and the coordinates of \mathbf{f} all have finite second moments. Define*

$$(\alpha, \beta, \gamma) = \underset{\alpha \in \mathbb{R}, \beta \in \mathbb{R}, \gamma \in \mathbb{R}^K}{\operatorname{argmin}} E[b(\alpha + u\beta + \mathbf{f}^T \gamma) - (\alpha + u\beta + \mathbf{f}^T \gamma)y].$$

We have the followings.

1. $|\beta| \geq |E(uy)|/(M \cdot E u^2)$;
2. If $E(uy) = 0$ and $P(u = 0) = 0$, then $\beta = 0$.

Now we investigate the sure screening property of the factor-adjusted screening procedure. Recall that $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_n)^T \in \mathbb{R}^{n \times K}$ and $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_n)^T \in \mathbb{R}^{n \times p}$ are matrices of true factors and idiosyncratic components, respectively, and their estimated versions are $\hat{\mathbf{F}}$ and $\hat{\mathbf{U}}$. We use $\mathbf{U}_j, \hat{\mathbf{U}}_j \in \mathbb{R}^n$ to refer to the j -th columns of \mathbf{U} and $\hat{\mathbf{U}}$. Define $L_{n,j}(\alpha, \beta, \gamma) = L_n(\mathbf{y}, \mathbf{1}_n \alpha + \hat{\mathbf{U}}_j \beta + \hat{\mathbf{F}} \gamma)$ for $j \in [p]$, $\alpha \in \mathbb{R}$, $\beta \in \mathbb{R}$ and $\gamma \in \mathbb{R}^K$. Let

$$(\alpha_j, \beta_j, \gamma_j) = \underset{\alpha \in \mathbb{R}, \beta \in \mathbb{R}, \gamma \in \mathbb{R}^K}{\operatorname{argmin}} E L_n(\mathbf{y}, \mathbf{1}_n \alpha + \mathbf{U}_j \beta + \mathbf{F} \gamma) \quad (4.2)$$

be the population version of $(\hat{\alpha}_j, \hat{\beta}_j, \hat{\gamma}_j)$.

The following three assumptions are variants of those in Section 4.1, and hence they hold almost surely or with high probability in the cases we are interested in.

Assumption 4.4 (Smoothness). $b \in C^3(\mathbb{R})$. For some constant M , we have $0 \leq b''(z) \leq M$ and $|b'''(z)| \leq M$, $\forall z$.

Assumption 4.5 (Strong convexity of marginal loss functions). There exists some constant $\kappa > 0$ such that $\nabla^2 L_{n,j}(\alpha_j, \beta_j, \gamma_j) \succeq \kappa \mathbf{I}$ for all $j \in [p]$.

Assumption 4.6 (Estimation of factor model). There exist constants C , c and a nonsingular matrix $\mathbf{H}_0 \in \mathbb{R}^{K \times K}$ such that the followings happen. Let $\mathbf{H} = \begin{pmatrix} \mathbf{I}_p & \mathbf{0}_{p \times K} \\ \mathbf{0}_{K \times p} & \mathbf{H}_0 \end{pmatrix}$ and $\overline{\mathbf{W}} = \hat{\mathbf{W}} \mathbf{H}$. We have $\|\mathbf{W}\|_{\max} \leq C$, $\|\overline{\mathbf{W}} - \mathbf{W}\|_{\max} \leq C$ and $\max_{j \in [p+K]} \left(\frac{1}{n} \sum_{t=1}^n |\bar{w}_{tj} - w_{tj}|^2 \right)^{1/2} \leq c$.

Theorem 4.2. Suppose that Assumptions 4.4, 4.5 and 4.6 hold with high probability, and the constant c in Assumption 4.6 is small enough. For $j \in [p-1]$, define

$$\epsilon_j = \frac{1}{n} \sum_{t=1}^n [b'(\alpha_j + \mathbf{f}_t^T \gamma_j + u_{tj} \beta_j) - y_t] (1, \mathbf{f}_t^T, u_{tj})^T.$$

If $\xi \leq \rho \min_{j \in \text{supp}(\beta^*)} |\mathbb{E}(u_j y)| / (M \cdot \mathbb{E} u_j^2)$ for some constant $\rho \in (0, 1)$ and

$$\max_{j \in \text{supp}(\beta^*)} \|\epsilon_j\|_2 = o_P(\min_{j \in \text{supp}(\beta^*)} |\mathbb{E}(u_j y)| / \mathbb{E} u_j^2), \quad (4.3)$$

then we have

$$\mathbb{P}(\text{supp}(\beta^*) \subseteq \{j : |\hat{\beta}_j| \geq \xi\}) = 1 - o(1).$$

Under the conditions in Lemma 4.1, we can prove that

$$\max_{j \in \text{supp}(\beta^*)} \|\epsilon_j\|_2 = O_P\left(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}}\right).$$

Theorem 4.2 asserts that if $\min_{j \in \text{supp}(\beta^*)} |\mathbb{E}(u_j y)| / \mathbb{E} u_j^2$ grows faster than $\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}}$, the factor-adjusted screening procedure enjoys the sure screening property (Fan and Lv, 2008; Fan and Song, 2009). The optimal choice of the screening threshold ξ can be discussed by following the analysis as in Fan and Song (2009). Here we do not pursue this result as it is not the main focus of the paper.

5 Simulation study

5.1 Example 1: Linear regression

We study a simulated example for high dimensional sparse linear regression with correlated covariates. The correlation structure is calibrated from S&P 500 monthly excess returns between 1980 and 2012. Throughout the numerical studies of this paper, the tuning parameter λ is selected by the 10-fold cross validation. The model selection performance is measured by the model selection consistency rate and the sure screening rate. The former is the proportion of simulations that the selected model is identical to the true one and the latter is the proportion of simulations that the selected model contains all important variables.

Calibration and data generation process

We calculate the centered monthly excess returns for the stocks in S&P 500 index that have complete records from January 1980 to December 2012. The data, collected from CRSP², contains the returns of 202 stocks with a time span of 396 months. Denote the centred monthly excess returns as \mathbf{z}_t , $t = 1, \dots, 396$. The calibration and data generation procedure are outlined as follows.

- (1) Fit \mathbf{z}_t with a three factor model. We apply PCA on the sample covariance of $\{\mathbf{z}_t\}_{t=1}^{396}$ and denote λ_k and $\boldsymbol{\xi}_k$, $k = 1, 2, 3$, as the top three eigenvalues and corresponding eigenvectors. We estimate loadings $\tilde{\mathbf{B}} = (\sqrt{\lambda_1}\boldsymbol{\xi}_1, \sqrt{\lambda_2}\boldsymbol{\xi}_2, \sqrt{\lambda_3}\boldsymbol{\xi}_3)$ and $\tilde{\mathbf{f}}_t = (\lambda_1^{-1/2}\boldsymbol{\xi}_1^T \mathbf{z}_t, \lambda_2^{-1/2}\boldsymbol{\xi}_2^T \mathbf{z}_t, \lambda_3^{-1/2}\boldsymbol{\xi}_3^T \mathbf{z}_t)^T$.
- (2) Calculate $\boldsymbol{\Sigma}_{\mathbf{B}}$ as the sample covariance of the rows of $\tilde{\mathbf{B}}$, which is $\text{diag}(\lambda_1, \lambda_2, \lambda_3)$. Generate loading matrix \mathbf{B} whose rows are draws from i.i.d. $N(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{B}})$.
- (3) Fit VAR(1) model $\tilde{\mathbf{f}}_t = \boldsymbol{\Phi}\tilde{\mathbf{f}}_{t-1} + \boldsymbol{\eta}_t$. Denote $\hat{\boldsymbol{\Phi}}$ the estimate of $\boldsymbol{\Phi}$ and calculate $\boldsymbol{\Sigma}_{\boldsymbol{\eta}} = \mathbf{I} - \hat{\boldsymbol{\Phi}}\hat{\boldsymbol{\Phi}}^T$. Generate \mathbf{f}_t from the VAR(1) model $\mathbf{f}_t = \hat{\boldsymbol{\Phi}}\mathbf{f}_{t-1} + \boldsymbol{\eta}_t$ with $\mathbf{f}_0 = \mathbf{0}$, where $\boldsymbol{\eta}_t$ is generated from i.i.d. $N(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\eta}})$.
- (4) Calculate the residual $\tilde{\mathbf{u}}_t = \mathbf{z}_t - \tilde{\mathbf{B}}\tilde{\mathbf{f}}_t$ and $\boldsymbol{\Sigma}_{\mathbf{u}}$ the sample covariance matrix of $\tilde{\mathbf{u}}_t$. Denote σ_u^2 the average of the diagonal entries of $\boldsymbol{\Sigma}_{\mathbf{u}}$. Generate covariates \mathbf{x}_t from a factor model $\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t$ where the entries in \mathbf{u}_t are drawn from i.i.d. $N(0, \sigma_u^2)$.
- (5) Generate the response y_t from a sparse linear model $y_t = \mathbf{x}_t^T \boldsymbol{\beta}^* + \varepsilon_t$. The true coefficients are $\boldsymbol{\beta}^* = (\beta_1, \dots, \beta_{10}, \mathbf{0}_{(p-10)}^T)^T$, and the nonzero coefficients $\beta_1, \dots, \beta_{10}$ are drawn from i.i.d. Uniform(2,5). We draw ε_t from an AR(1) model $\varepsilon_t = 0.5\varepsilon_{t-1} + \gamma_t$ with $\gamma_t \sim N(0, 0.3)$.

The results of the calibrated parameters are presented in Table 1.

Table 1: Parameters calibrated from S&P 500 returns

$\boldsymbol{\Sigma}_{\mathbf{B}}$			$\hat{\boldsymbol{\Phi}}$			$\boldsymbol{\Sigma}_{\boldsymbol{\eta}}$			σ_u^2
0.5237	0	0	0.1897	-0.0375	-0.0223	0.9621	-0.0056	0.0182	0.0146
0	0.2884	0	0.0630	0.1553	0.0206	-0.0056	0.9715	-0.0078	
0	0	0.2372	-0.0432	0.0102	0.4343	0.0182	-0.0078	0.8094	

²Center for Research in Security Prices Database, see <http://www.crsp.com/> for more details.

Impacts of Irrepresentable Condition

First, we show LASSO performs poorly in terms of model selection consistency rate when the *irrepresentable condition* is violated, while FarmSelect can consistently select the correct model. Let $n = 100$ and $p = 500$. Denote by $\Gamma_\infty = \|\mathbf{X}_{S^c}^T \mathbf{X}_S (\mathbf{X}_S^T \mathbf{X}_S)^{-1}\|_\infty$. When $\Gamma_\infty < 1$ the *irrepresentable condition* holds and otherwise it is violated. We simulate 10,000 replications. For each replication, we calculate Γ_∞ and apply both LASSO and FarmSelect for model selection. Then we calculate the model selection consistency rate within each small interval around Γ_∞ (a nonparametric smoothing). The results are presented in Figure 2. According to Figure 2, both FarmSelect and LASSO have high model selection consistency rate when $\Gamma_\infty < 1$. This shows FarmSelect does not pay any price under the weak correlation scenario. As Γ_∞ grows beyond 1, the correct model selection rate of LASSO drops quickly. When the *irrepresentable condition* is strongly violated (e.g. $\Gamma_\infty > 1.5$), the correct model selection rate of LASSO is close to zero. On the contrary, FarmSelect has high selection consistency rates regardless of Γ_∞ .




Figure 2: Relationship between model selection consistency rate and *irrepresentable condition*. Among the 10,000 replications, more than 9,500 replications have $\Gamma_\infty > 1$ and more than 8,000 replications have $\Gamma_\infty > 1.5$.

Impacts of sample size

Second, we examine the model selection consistency with a fixed dimensionality and an increasing sample size. We fix $p = 500$ and let n increase from 50 to 150. For each given sample size, we simulate 200 replications and calculate the model selection consistency rates and the sure screening rates for LASSO, SCAD, elastic net and FarmSelect, respectively. For the elastic net, we set $\lambda_1 = \lambda_2 \equiv \lambda$. The results are presented in Figure 3 (a) and Figure 3 (b). Figure 3 (a) shows that model selection consistency rates of LASSO, SCAD and elastic net do not enjoy fast convergence to 1 when the sample size increases, while the one of FarmSelect equals to one as long as the sample size exceeds 100. Similar phenomena are observed from sure screening rates. To demonstrate the prediction performance, we report the mean estimation error $\|\hat{\beta} - \beta^*\|_2$ for each method, which is a good indicator of the prediction error. The estimation errors are reported in Figure 3 (c). When the sample size is small, LASSO, SCAD and elastic net have large estimation errors since they tend to select overfitted models.

Impacts of dimensionality

Third, we assess the model selection performance when the dimensionality p is growing beyond n and diverging. We fix $n = 100$ and let p grow from 200 to 1000. For each given p , we simulate 200 replications and calculate the model selection consistency rate of LASSO, SCAD, elastic net and FarmSelect respectively. The model selection consistency rates are presented in Figure 4(a). According to Figure 4(a), the model selection consistency rate of FarmSelect stays close to 1 even as p increases, whereas the rates for the other three methods drop quickly. Again, we report the estimation errors in Figure 4(b). As the dimensionality grows, FarmSelect has the least increase of estimation error.

5.2 Example 2: Logistic regression

We consider the following logistic regression model whose conditional probability function is:

$$P(y_t = 1 | \mathbf{X}_t) = \frac{\exp(\mathbf{X}_t^T \beta)}{1 + \exp(\mathbf{X}_t^T \beta)}, \quad i = 1, \dots, N. \quad (5.1)$$

We set sample size $n = 200$ and dimensionality $p = 300, 400, 500$. The true coefficients are set to be $\beta^* = (\beta_{(1)}^T, \mathbf{0})^T$ with $\beta_{(1)} = (6, 5, 4)^T$. Hence the true model size is 3.

The covariates \mathbf{X} are generated from one of the following three models:




Figure 3: From above to below: (a) Model selection consistency rates with fixed p and increasing n ; (b) Sure screening rates with fixed p and increasing n ; (c) Mean estimation errors $\|\hat{\beta} - \beta^*\|_2$ with fixed p and increasing n .




Figure 4: From above to below: (a) Model selection consistency rates with fixed n and increasing p ; (b) Mean estimation errors $\|\hat{\beta} - \beta^*\|_2$ with fixed n and increasing p .

- (1) Factor model $\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t$ with $K = 3$. Factors are generated from a stationary VAR(1) model $\mathbf{f}_t = \Phi\mathbf{f}_{t-1} + \boldsymbol{\eta}_t$ with $\mathbf{f}_0 = \mathbf{0}$. The (i, j) th entry of Φ is set to be 0.5 when $i = j$ and $0.3^{|i-j|}$ when $i \neq j$. We draw \mathbf{B} , \mathbf{u}_t and $\boldsymbol{\eta}_t$ from the i.i.d. standard Normal distribution.
- (2) Equal correlated case. We draw \mathbf{x}_t from i.i.d. $N_p(\mathbf{0}, \Sigma)$, where Σ has diagonal elements 1 and off-diagonal elements 0.8.
- (3) Uncorrelated case. We draw \mathbf{x}_t from i.i.d. $N_p(\mathbf{0}, \mathbf{I})$

We compare the model selection performance of FarmSelect with LASSO and simulate 100

replications for each scenario. The model selection performance is measured by selection consistency rate, sure screening rate and the average size of selected model. The results are presented in Table 2 below. According to Table 2, FarmSelect pays no price for the uncorrelated case and outperforms LASSO for highly correlated cases.

Table 2: Model selection results of logistic regression ($n = 200$)

	Factor model with $K = 3$					
	FarmSelect			LASSO		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.94	1.00	3.07	0.07	0.98	12.61
$p = 400$	0.90	0.99	3.12	0.05	0.95	12.94
$p = 500$	0.83	0.98	3.18	0.03	0.93	15.07
	Equal correlated case					
	FarmSelect			LASSO		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.93	1.00	3.09	0.07	0.85	9.90
$p = 400$	0.89	1.00	3.14	0.05	0.80	10.82
$p = 500$	0.85	0.99	3.19	0.02	0.69	11.79
	Uncorrelated case					
	FarmSelect			LASSO		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.97	1.00	3.03	0.95	1.00	3.14
$p = 400$	0.93	1.00	3.07	0.91	1.00	3.34
$p = 500$	0.91	1.00	3.10	0.89	1.00	3.42

6 Prediction of U.S. bond risk premia

In this section, we predict U.S. bond risk premia with a large panel of macroeconomic variables. The response variable is the monthly data of U.S. bond risk premia with maturity of 2 to 5 years

between January 1980 and December 2015 containing 432 data points. The bond risk premia is calculated as the one year return of an n years maturity bond excessing the one year maturity bond yield as the risk-free rate. The covariates are 134 monthly U.S. macroeconomic variables in the FRED-MD database³ (McCracken and Ng, 2016). The covariates in the FRED-MD dataset are strongly correlated and can be well explained by a few principal components. To see this, we apply principal component analysis to the covariates and draw the scree plot of the top 20 principal components in Figure 5. The scree plot shows the first principal component solely explains more than 60% of the total variance. In addition, the first 5 principal components together explain more than 90% of the total variance.

We apply one month ahead rolling window prediction with a window size of 120 months. Within each window, we predict the U.S. bond risk premia by a high dimensional linear regression model of dimensionality 134. We compare the proposed FarmSelect method with LASSO in terms of model selection and prediction. In addition, we include the principal component regression (PCR) in the competition of prediction. Instead of using the covariates as regressors directly, PCR regresses the dependent variable on the leading principal components of covariates. The FarmSelect is implemented by the *FarmSelect* R package with default settings. To be specific, the loss function is L_1 , the number of factors is estimated by the modified eigen-ratio method and the regularized parameter is selected by multi-fold cross validation. The LASSO method is implemented by the *glmnet* R package. The PCR method is implemented by the *pls* package in R. The number of principal components is chosen as 8 which is suggested in Ludvigson and Ng (2009).

The prediction performance is evaluated by the out-of-sample R^2 which is defined as

$$R^2 = 1 - \frac{\sum_{t=121}^{432} (y_t - \hat{y}_t)^2}{\sum_{t=121}^{432} (y_t - \bar{y}_t)^2},$$

where y_t is the response variable realized at time t , \hat{y}_t is the predicted y_t by one of the three methods above using the previous 120 months data, and \bar{y}_t is the sample mean of the previous 120 months responses $(y_{t-120}, \dots, y_{t-1})$, which represents a naive predictor. For FarmSelect and LASSO, we also report the average selected model size for prediction at time $t \in \{121, \dots, 432\}$. The out-of-sample R^2 and average selected model size are reported in Table 3. The detailed prediction performances can be viewed in Figure 6. The results in Table 3 show that FarmSelect selects parsimonious models and achieves the highest R^2 's in all scenarios. On the contrary, LASSO may select redundant models

³The FRED-MD is a monthly economic database updated by the Federal Reserve Bank of St. Louis which is public available at <http://research.stlouisfed.org/econ/mccracken/sel/>.

as it ignores the correlations among covariates. To see this, we rank the covariates according to the selection frequency. The top 10 selected covariates and their frequencies are listed in Table 4. According to Table 4, LASSO tends to select some highly correlated covariates simultaneously. For instance, LASSO includes both Housing Starts Northeast and New Private Housing Permits, Northeast (SAAR) due to strong correlation between them. In addition, both Switzerland/U.S. and Japan/U.S. exchange rates enter the solution path of LASSO early, which can be another evidence of overfitting.




Figure 5: Eigenvalues (dotted line) and proportion of variance explained (bar) by the top 20 principal components




Figure 6: One month ahead rolling window forecast. The window size is 120 months. The solid black line is the true bond risk premia. From top to bottom are the plots for 2 to 5 years bond risk premia

Table 3: Out of sample R^2 and average selected model size

Maturity of Bond	Out of sample R^2			Average model size	
	FarmSelect	LASSO	PCR	FarmSelect	Lasso
2 Years	0.543	0.507	0.434	9.04	22.72
3 Years	0.512	0.461	0.419	9.52	21.40
4 Years	0.475	0.412	0.382	9.80	20.74
5 Years	0.440	0.387	0.355	10.07	20.21

References

- AKAIKE, H. (1973). Information theory and an extension of the maximum likelihood principle. In *Second International Symposium in Information Theory*, (B.N. Petroc and F. Caski, eds.). Akademiai Kiado, Budapest, 276–281.
- AHN, S. C. and HORENSTEIN, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, **81**, 1203–1227.
- BAI, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, **71**, 135–171.
- MONTGOMERY, A. L., ZARNOWITZ, V., TSAY, R. S. and TIAO, G. C. (1998). Forecasting the US unemployment rate. *Journal of the American Statistical Association*, **93**(442), 478–493.
- LIU, L. M., HUDAK, G. B., BOX, G. E., MULLER, M. E. and TIAO, G. C. (1992). *Forecasting and time series analysis using the SCA statistical system*, Vol 1(2), DeKalb, IL: Scientific Computing Associates.
- BOX, G. E. and TIAO, G. C. (1976). Comparison of forecast and actuality. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **25**(3), 195–200.
- BAI, J. and NG, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, **70**, 191–221.
- BICKEL, P. J., RITOV, Y. A. and TSYBAKOV, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *The Annals of statistics*, **37**, 1705–1732.

Table 4: 2 years Maturity: Top 10 variables with highest selection frequency

FarmSelect		
Rank	Name	Frequency
1	Switzerland / U.S. Foreign Exchange Rate	79
2	Civilians Unemployed - Less Than 5 Weeks	76
3	Moody's Baa Corporate Bond Minus FEDFUNDS	74
4	Housing Starts, Northeast	73
5	Industrial Production: Durable Consumer Goods	70
6	CBOE S&P 100 Volatility Index	68
7	Real M2 Money Stock	67
8	10-Year Treasury Rate	65
9	CPI : Commodities	62
10	Commercial and Industrial Loans	60
LASSO		
Rank	Name	Frequency
1	CBOE S&P 100 Volatility Index	134
2	Industrial Production: Residential Utilities	130
3	Housing Starts, Northeast	127
4	Switzerland / U.S. Foreign Exchange Rate	126
5	Industrial Production: Fuels	124
6	New Private Housing Permits, South	122
7	Canada / U.S. Foreign Exchange Rate	117
8	10-Year Treasury Rate	114
9	Japan / U.S. Foreign Exchange Rate	110
10	CPI : Commodities	106

BOYD, S., PARIKH, N., CHU, E., PELEATO, B., and ECKSTEIN, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends[®] in Machine Learning*, **3**, 1–122.

- CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, **35**, 2313–2351.
- CHANG, J., GUO, B. and YAO, Q. (2015). High dimensional stochastic regression with latent factors, endogeneity and nonlinearity. *Journal of Econometrics*, **189**, 297–312.
- CHAMBERLAIN, G. and ROTHCHILD, M. (1982). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, **51**, 1305–1324.
- CHOI, B.S. (1992). *ARMA Model Identification.*, Springer-Verlag, New York.
- DONOHOO, D. L. and ELAD, M. (2003). Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization. *Proceedings of the National Academy of Sciences*, **100**, 2197–2202.
- EFRON, B., HASTIE, T., JOHNSTONE, I. and TIBSHIRANI, R. (2004). Least angle regression. *The Annals of statistics*, **32**, 407–499.
- FAN, J., and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, **96**, 1348–1360.
- FAN, J., LIAO, Y. and MINCHEVA, M. (2013). Large covariance estimation by thresholding principal orthogonal complements (with discussion). *Journal of the Royal Statistical Society, Series B*, **75**, 603–680.
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society, Series B*, **70**, 849–911.
- FAN, J. and PENG, H. (2004). On non-concave penalized likelihood with diverging number of parameters. *Annals of Statistics*, **32**, 928–961.
- FAN, J. and SONG, R. (2009). Sure Independence Screening in Generalized Linear Models with NP-Dimensionality. *Annals of Statistics*, **38**, 3567–3604.
- FAN, J., WANG, K., ZHONG, Y., and ZHU, Z. (2019+). Robust high dimensional factor models with applications to statistical machine learning. *Statistical Science*, invited.
- FORNI, M., HALLIN, M., LIPPI, M. and REICHLIN, L. (2005). The generalized dynamic factor model: one-sided estimation and forecasting. *Journal of the American Statistical Association*, **100**, 830–840.

- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, **33**, 1–22.
- HALLIN, M. and LIŠKA, R. (2007). Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association*, **102**, 603–617.
- JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, **104**, 682–693.
- KNEIP, A. and SARDA, P. (2011). Factor models and variable selection in high-dimensional regression analysis. *The Annals of Statistics*, **39**, 2410–2447.
- LAM, C. and YAO, Q. (2012). Factor modeling for high dimensional time-series: inference for the number of factors. *Annals of Statistics*, **40**, 694–726.
- LAWLEY, D. and MAXWELL, A. (1971). *Factor analysis as a statistical method.*, Butterworths, London.
- LEE, J. D., SUN, Y. and TAYLOR, J. E. (2015). On model selection consistency of regularized M-estimators. *Electronic Journal of Statistics*, **9**, 608–642.
- LUO, R., WANG, H. and TSAI, C. L. (2009). Contour projected dimension reduction. *Ann. Statist.*, **37**, 3743–3778.
- LUDVIGSON, S. and NG, S. (2009). Macro factors in bond risk premia. *Review of Financial Studies*, **22**, 5027–5067.
- MCCRACKEN, M. and NG, S. (2016). FRED-MD: A Monthly Database for Macroeconomic Research. *Journal of Business & Economic Statistics*, **34**, 574–589.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, **34**, 1436–1462.
- MERLEVÈDE, FLORENCE and PELIGRAD, MAGDA and RIO, EMMANUEL (2011). A Bernstein type inequality and moderate deviations for weakly dependent sequences. *Probability Theory and Related Fields*, **151**, 435–474.
- NEGAHBAN, S., RAVIKUMAR, P. K., WAINWRIGHT, M. J. and YU, B. (2012). A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. *Statistical Science*, **27**, 1348–1356.

- OBERTHUER, A., BERTHOLD, F., WARNAT, P., HERO, B., KAHLERT, Y., SPITZ, R., ERNESTUS, K., KÖNIG, R., HAAS, S., EILS, R., and SCHWAB, M. (2006). Customized oligonucleotide microarray gene expression-based classification of neuroblastoma patients outperforms current clinical risk stratification. *Journal of clinical oncology*, **24**, 5070–5078.
- SCHWARZ, G.E. (1978). Estimating the dimension of a model. *The Annals of statistics*, **6**, 461–464.
- SHEN, D., SHEN, H., ZHU, H. and MARRON, J. S. (2016). The statistics and mathematics of high dimension low sample size asymptotics. *Statistica Sinica*, **26**, 1747–1770.
- STOCK, J. and WATSON, M. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, **97**, 1167–1179.
- TIAO, G. C., and TSAY, R. S. (1989). Model specification in multivariate time series. *Journal of the Royal Statistical Society. Series B*, **51**, 157–213.
- TSAY, R. S., and TIAO, G. C. (1985). Use of canonical analysis in time series model identification. *Biometrika*, **72**, 299–315.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B*, **58**, 267–288.
- VAN DE GEER, S. A. (2008). High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, **36**, 614–645.
- WAINWRIGHT, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso). quadratic programming (Lasso) *IEEE transactions on information theory*, **55**, 2183–2202.
- WANG, H (2012). Factor profiled sure independence screening. *Biometrika*, **99**, 15–28.
- WANG, W. and FAN, J. (2017). Asymptotics of Empirical Eigen-structure for Ultra-high Dimensional Spiked Covariance Model. *Annals of Statistics*, **45**, 1342–1374.
- WANG, X. and LENG, C. (2016). High dimensional ordinary least squares projection for screening variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **78**, 589–611.
- ZHANG, C. H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, **38**, 894–942.

- ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *Journal of Machine Learning Research*, **7**, 2541–2563.
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **67**, 301–320.
- ZOU, H. and LI, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of statistics*, **36**, 1509–1533.
- LAWLEY, D. and MAXWELL, A. (1971). *Factor analysis as a statistical method.*, Butterworths, London.

A Some preliminary results

In the first appendix, we introduce some useful results in convex analysis and inverse problems. Under mild conditions, the tools we developed connect the unique global optimum of the regularized loss function $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ with the solution of a constrained problem $\min_{\text{supp}(\boldsymbol{\theta}) \subseteq \mathcal{S}} L_\lambda(\boldsymbol{\theta})$.

Lemma A.1. *Suppose $L(\boldsymbol{\theta}) \in C^2(\mathbb{R}^p)$ and is convex. $R(\boldsymbol{\theta})$ is convex and $R(\boldsymbol{\alpha} + \boldsymbol{\beta}) = R(\boldsymbol{\alpha}) + R(\boldsymbol{\beta})$ for $\boldsymbol{\alpha} \in \mathcal{M}$ and $\boldsymbol{\beta} \in \mathcal{M}^\perp$, where \mathcal{M} is a linear subspace of \mathbb{R}^p and \mathcal{M}^\perp is its orthonormal complement. In addition, there exists $R^*(\boldsymbol{\theta}) \in C(\mathbb{R}^p)$ such that $|\langle \boldsymbol{\alpha}, \boldsymbol{\beta} \rangle| \leq R(\boldsymbol{\alpha})R^*(\boldsymbol{\beta})$ for $\boldsymbol{\alpha} \in \mathcal{M}^\perp$ and $\boldsymbol{\beta} \in \mathbb{R}^p$. Let $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ where $\lambda \geq 0$, and $\hat{\boldsymbol{\theta}} \in \text{argmin}_{\boldsymbol{\theta} \in \mathcal{M}} L_\lambda(\boldsymbol{\theta})$.*

If $R^(\nabla L(\hat{\boldsymbol{\theta}})) < \lambda$ and $\boldsymbol{\theta}^T \nabla^2 L(\hat{\boldsymbol{\theta}}) \boldsymbol{\theta} > 0$ for all $\boldsymbol{\theta} \in \mathcal{M}$, then $\hat{\boldsymbol{\theta}}$ is the unique global minimizer of $L_\lambda(\boldsymbol{\theta})$.*

Proof. For any $\boldsymbol{\theta} \in \mathbb{R}^p$ we use $\boldsymbol{\theta}_\mathcal{M}$, $\boldsymbol{\theta}_{\mathcal{M}^\perp}$ to denote its orthonormal projections on \mathcal{M} and \mathcal{M}^\perp , respectively. On the one hand, by convexity and orthogonality we have

$$L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M}) \geq \langle \nabla L(\boldsymbol{\theta}_\mathcal{M}), \boldsymbol{\theta} - \boldsymbol{\theta}_\mathcal{M} \rangle = \langle \nabla L(\boldsymbol{\theta}_\mathcal{M}), \boldsymbol{\theta}_{\mathcal{M}^\perp} \rangle \geq -R(\boldsymbol{\theta}_{\mathcal{M}^\perp})R^*(\nabla L(\boldsymbol{\theta}_\mathcal{M})).$$

Since $R^*(\nabla L(\hat{\boldsymbol{\theta}})) < \lambda$, there exists $\delta > 0$ such that $\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2 < \delta$ implies $R^*(\nabla L(\boldsymbol{\theta})) < \lambda$. Together with $\|\boldsymbol{\theta}_\mathcal{M} - \hat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2$, we know $L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M}) \geq -\lambda R(\boldsymbol{\theta}_{\mathcal{M}^\perp})$ as long as $\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2 < \delta$, and the inequality strictly holds when $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$.

On the other hand, $R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}_\mathcal{M}) = R(\boldsymbol{\theta}_\mathcal{M} + \boldsymbol{\theta}_{\mathcal{M}^\perp}) - R(\boldsymbol{\theta}_\mathcal{M}) = R(\boldsymbol{\theta}_{\mathcal{M}^\perp})$. Hence $\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2 < \delta$ forces $L_\lambda(\boldsymbol{\theta}) - L_\lambda(\boldsymbol{\theta}_\mathcal{M}) = [L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M})] + \lambda[R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}_\mathcal{M})] \geq 0$ and the inequality strictly holds when $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$.

Now suppose $0 < \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2 < \delta$. If $\boldsymbol{\theta} \in \mathcal{M}$, then the facts $\hat{\boldsymbol{\theta}} \in \text{argmin}_{\boldsymbol{\theta} \in \mathcal{M}} L_\lambda(\boldsymbol{\theta})$ and $\boldsymbol{\alpha}^T \nabla^2 L(\hat{\boldsymbol{\theta}}) \boldsymbol{\alpha} > 0$, $\forall \boldsymbol{\alpha} \in \mathcal{M}$ implies that $L_\lambda(\boldsymbol{\theta}') > L_\lambda(\hat{\boldsymbol{\theta}})$. In addition, our assumptions yield $\|\boldsymbol{\theta}\|_2^2 \leq R(\boldsymbol{\theta})R^*(\boldsymbol{\theta})$ for $\boldsymbol{\theta} \in \mathcal{M}^\perp$, leading to $R(\boldsymbol{\theta}) > 0$ over $\mathcal{M}^\perp \setminus \{0\}$. If $\boldsymbol{\theta} \notin \mathcal{M}$, then $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$ and $L_\lambda(\boldsymbol{\theta}) > L_\lambda(\boldsymbol{\theta}_\mathcal{M}) \geq L_\lambda(\hat{\boldsymbol{\theta}})$. Therefore $\hat{\boldsymbol{\theta}}$ is a strict local optimum of $L_\lambda(\boldsymbol{\theta})$, which is convex over \mathbb{R}^p . This finishes the proof. □

Lemma A.2. *Let $L(\boldsymbol{\theta})$ be convex over a Euclidean space \mathcal{M} . If $\boldsymbol{\theta}_0 \in \mathcal{M}$, $r > 0$, and $L(\boldsymbol{\theta}) > L(\boldsymbol{\theta}_0)$ over the sphere $\partial B(\boldsymbol{\theta}_0, r)$, then any minimizer of $L(\boldsymbol{\theta})$ is within the ball $B(\boldsymbol{\theta}_0, r)$.*

Proof. For any $\boldsymbol{\theta} \notin \overline{B(\boldsymbol{\theta}_0, r)}$, there exists $t \in (0, 1)$ and $\boldsymbol{\theta}' \in \partial B(\boldsymbol{\theta}_0, r)$ such that $\boldsymbol{\theta}' = (1-t)\boldsymbol{\theta} + t\boldsymbol{\theta}_0$. Then $L(\boldsymbol{\theta}_0) < L(\boldsymbol{\theta}') \leq (1-t)L(\boldsymbol{\theta}) + tL(\boldsymbol{\theta}_0)$, yielding $L(\boldsymbol{\theta}) > L(\boldsymbol{\theta}_0)$. Hence there is no minimizer outside $B(\boldsymbol{\theta}_0, r)$. \square

Lemma A.3. Suppose \mathcal{M} is a Euclidean space, $\boldsymbol{\theta}_0 \in \mathcal{M}$ and $L(\boldsymbol{\theta})$ is convex over \mathcal{M} . In addition, there exist $\kappa, A > 0$ such that $L(\boldsymbol{\theta}) \geq L(\boldsymbol{\theta}_0) + \langle \mathbf{h}, \boldsymbol{\theta} - \boldsymbol{\theta}_0 \rangle + \frac{\kappa}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2$ as long as $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$ and $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq A$. If $\inf_{\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$, then any minimizer of $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ is within the ball $\{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa} \inf_{\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2\}$.

Proof. If $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 < A$ and $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$, then

$$\begin{aligned} L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_0) &\geq \langle \mathbf{h}, \boldsymbol{\theta} - \boldsymbol{\theta}_0 \rangle + \frac{\kappa}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2 \geq -\|\mathbf{h}\|_2 \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 + \frac{\kappa}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2 \\ &= \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \left(\frac{\kappa}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 - \|\mathbf{h}\|_2 \right). \end{aligned}$$

Taking $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$ and $r > 0$ such that $\frac{2}{\kappa} \|\mathbf{h}\|_2 < r < A$. This forces $L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_0) > 0$ over the sphere $B(\boldsymbol{\theta}_0, r)$. Let $\hat{\boldsymbol{\theta}}$ be one of the minimizers of $L(\boldsymbol{\theta})$. Lemma A.2 implies that $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\| < r < A$. Then $0 \geq L(\hat{\boldsymbol{\theta}}) - L(\boldsymbol{\theta}_0) \geq \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \left(\frac{\kappa}{2} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 - \|\mathbf{h}\|_2 \right)$. The result is proved by taking the infimum over $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$. \square

Corollary A.1. Suppose $\lambda \geq 0$, \mathcal{M} is a Euclidean space, $\boldsymbol{\theta}_0 \in \mathcal{M}$, $L(\boldsymbol{\theta}) \in C^2(\mathcal{M})$ and is convex, and $R(\boldsymbol{\theta})$ is convex. In addition, there exist $\kappa, A > 0$ such that $\nabla^2 L(\boldsymbol{\theta}) \succeq \kappa I$ as long as $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq A$. If $\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \inf_{\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$, then $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ has unique minimizer $\hat{\boldsymbol{\theta}}$ and $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa} (\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \inf_{\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2)$.

Proof. Note that $\partial L_\lambda(\boldsymbol{\theta}_0) = \nabla L(\boldsymbol{\theta}_0) + \lambda \partial R(\boldsymbol{\theta}_0)$. There exists $\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)$ such that $\mathbf{h}' = \nabla L(\boldsymbol{\theta}_0) + \mathbf{h} \in \partial L_\lambda(\boldsymbol{\theta}_0)$ and $\|\mathbf{h}'\|_2 < \|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$. Applying Lemma A.3 to L_λ and \mathbf{h}' , we obtain that any minimizer of L_λ satisfies $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa} \|\mathbf{h}'\|_2 < \frac{2}{\kappa} (\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \|\mathbf{h}\|_2)$. Then $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq A$ and $\nabla^2 L(\hat{\boldsymbol{\theta}}) \succ 0$, proving both the bound and uniqueness. \square

Proof of Lemma 3.1

Let $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_n)^T$ and $\boldsymbol{\theta}^* = ((\boldsymbol{\beta}^*)^T, (\boldsymbol{\beta}^*)^T \mathbf{B}_0)^T$. Note that $\nabla E[L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta})] = E\{\frac{1}{n} \sum_{t=1}^n [-y_t + b(\mathbf{w}_t^T \boldsymbol{\theta})] \mathbf{w}_t\} = E\{-y_1 + b(\mathbf{w}_1^T \boldsymbol{\theta})\} \mathbf{w}_1$ and $\mathbf{w}_t^T \boldsymbol{\theta}^* = (1, \mathbf{x}_t^T) \boldsymbol{\beta}^*$. The claim is proved by

$$\nabla E[L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta})]_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} = E\{-y_1 + b(\mathbf{w}_1^T \boldsymbol{\theta}^*)\} \mathbf{w}_1 = E\{-y_1 + b((1, \mathbf{x}_1^T) \boldsymbol{\beta}^*)\} \mathbf{w}_1 = E(\eta_1 \mathbf{w}_1) = \mathbf{0}.$$

Proof of Lemma 3.2

Lemma 3.2 is similar to the results developed in the Appendix C of Wang and Fan (2017) and hence we omit the details.

B General results for M -estimators

We present general model selection results to be used as tools later. Without loss of generality, we assume the last K variables are not penalized. Let $L_n : \mathbb{R}^{p+K} \rightarrow \mathbb{R}$ be a convex loss function, $\boldsymbol{\theta}^* \in \mathbb{R}^{p+K}$ and $\boldsymbol{\beta}^* = \boldsymbol{\theta}_{[p]}^*$ be the sparse sub-vector of interest. Then $\boldsymbol{\theta}^*$ and $\boldsymbol{\beta}^*$ are estimated via

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \mathbb{R}^{p+K}}{\operatorname{argmin}} \{L_n(\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[p]}\|_1\} \quad \text{and} \quad \hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\theta}}_{[p]},$$

respectively. Further, denote $S = \operatorname{supp}(\boldsymbol{\theta}^*)$, $S_1 = \operatorname{supp}(\boldsymbol{\beta}^*)$ and $S_2 = [p+K] \setminus S$.

Assumption B.1 (Smoothness). $L_n(\boldsymbol{\theta}) \in C^2(\mathbb{R}^{p+K})$ and there exist $A > 0$, $M > 0$ such that $\|\nabla_{S'}^2 L(\boldsymbol{\theta}) - \nabla_{S'}^2 L_n(\boldsymbol{\theta}^*)\|_\infty \leq M \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2$ whenever $\operatorname{supp}(\boldsymbol{\theta}) \subseteq S$ and $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq A$.

Assumption B.2 (Restricted strong convexity). There exist $\kappa_2 > \kappa_\infty > 0$ such that $\|[\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)]^{-1}\|_\infty \leq \frac{1}{2\kappa_\infty}$ and $\|[\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)]^{-1}\|_2 \leq \frac{1}{2\kappa_2}$.

Assumption B.3 (Irrepresentable condition). $\|\nabla_{S_2S}^2 L_n(\boldsymbol{\theta}^*)[\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)]^{-1}\|_\infty \leq 1 - \tau$ for some $\tau \in (0, 1)$.

Assumptions B.1 – B.3 are standard in the studies of high-dimensional regularized estimators (e.g. Negahban et al., 2012; Lee et al., 2015). Based on them, we introduce the following theorem of L^p ($p = 1, 2, \infty$) error bounds and sign consistency for the FarmSelect estimator.

Theorem B.1. (i) **Error bounds** : Under Assumptions B.1 – B.3, if

$$\frac{7}{\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty < \lambda < \frac{\kappa_2}{4\sqrt{|S|}} \min \left\{ A, \frac{\kappa_\infty \tau}{3M} \right\}, \quad (\text{B.1})$$

then $\operatorname{supp}(\hat{\boldsymbol{\theta}}) \subseteq S$ and

$$\begin{aligned} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty &\leq \frac{3}{5\kappa_\infty} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda), \\ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 &\leq \frac{2}{\kappa_2} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|}), \\ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 &\leq \min \left\{ \frac{3}{5\kappa_\infty} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \lambda|S_1|), \frac{2\sqrt{|S|}}{\kappa_2} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|}) \right\}. \end{aligned}$$

(ii) **Sign consistency** : In addition, if the following two conditions

$$\begin{aligned} \min\{|\beta_j^*| : \beta_j^* \neq 0, j \in [p]\} &> \frac{C}{\kappa_\infty \tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty, \\ \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty &< \frac{\kappa_2 \tau}{7C\sqrt{|S|}} \min\left\{A, \frac{\kappa_\infty \tau}{3M}\right\} \end{aligned} \quad (\text{B.2})$$

hold for some $C \geq 5$, then by taking $\lambda \in (\frac{7}{\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty, \frac{1}{\tau}(\frac{5C}{3} - 1) \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty)$, the estimator achieves the sign consistency $\text{sign}(\hat{\boldsymbol{\beta}}) = \text{sign}(\boldsymbol{\beta}^*)$.

Remark B.1. Theorem B.1 shows how the correlated covariates affect the sign consistency as well as error bounds. To achieve the sign consistency, the tuning parameter λ should scale with τ^{-1} . Therefore, the L_∞ and L_2 errors will scale with $(\kappa_\infty \tau)^{-1}$ and $(\kappa_2 \tau)^{-1}$, respectively. When the covariates are highly correlated, the *irrepresentable condition* will get violated or the parameter $\tau \in (0, 1)$ is very small. As a result, the model selection procedures will fail to achieve the sign consistency and the error bounds will be suboptimal. On the other hand, the optimal error bounds require a small λ , which typically leads to an overfitted model. One can see a trade-off between model selection and parameter estimation due to the existence of dependency.

Remark B.2. The L_1 and L_2 error bounds in Theorem B.1 depend on $|S_1|$, the number of active variables. They stem from the bias induced by the penalty. To reduce the bias, it is desirable to penalize as few active variables as possible. This phenomena motivates FarmSelect to adopt a penalized profile likelihood form by not imposing penalty on the nuisance parameter γ .

As discussed in Remark B.1, when the covariates are highly correlated, the *irrepresentable condition* may not hold, or has a very small τ . This makes the model selection consistency either very hard to achieve or incompatible with low estimation error bounds. Therefore, the FarmSelect strategy can improve the model selection consistency and reduce the estimation error bounds if \mathbf{X} can be decomposed into $\mathbf{F}\mathbf{B}^T + \mathbf{U}$ such that (\mathbf{U}, \mathbf{F}) is well-behaved. This is due to the fact that the *irrepresentable condition* is easier to hold with positive τ bounded away from zero after the decorrelation step. To this end, any effective decorrelation procedure can be incorporated into this frame work.

B.1 Proof of Theorem B.1

Define $B_S(\boldsymbol{\theta}^*, r) = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq r, \text{supp}(\boldsymbol{\theta}) \subseteq S\}$ for $r > 0$. We first introduce two useful lemmas.

Lemma B.1. Suppose $\mathbf{A} \in \mathbb{R}^{q \times r}$ and $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{r \times r}$ and $\|\mathbf{CB}^{-1}\| < 1$, where $\|\cdot\|$ is an induced norm. Then $\|\mathbf{A}[(\mathbf{B} + \mathbf{C})^{-1} - \mathbf{B}^{-1}]\| \leq \frac{\|\mathbf{AB}^{-1}\| \cdot \|\mathbf{CB}^{-1}\|}{1 - \|\mathbf{CB}^{-1}\|}$.

Proof. By the sub-multiplicity of induced norms,

$$\begin{aligned} \|\mathbf{A}[(\mathbf{B} + \mathbf{C})^{-1} - \mathbf{B}^{-1}]\| &= \|\mathbf{AB}^{-1}[(\mathbf{I} + \mathbf{CB}^{-1})^{-1} - \mathbf{I}]\| \leq \|\mathbf{AB}^{-1}\| \cdot \|(\mathbf{I} + \mathbf{CB}^{-1})^{-1} - \mathbf{I}\| \\ &= \|\mathbf{AB}^{-1}\| \cdot \left\| \sum_{k=0}^{\infty} (-\mathbf{CB}^{-1})^k - \mathbf{I} \right\| \leq \|\mathbf{AB}^{-1}\| \sum_{k=1}^{\infty} \|\mathbf{CB}^{-1}\|^k = \frac{\|\mathbf{AB}^{-1}\| \cdot \|\mathbf{CB}^{-1}\|}{1 - \|\mathbf{CB}^{-1}\|}. \end{aligned}$$

□

Lemma B.2. Under Assumptions B.1 and B.2, we have $\|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1}\|_2 < \kappa_2^{-1}$ and $\|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1}\|_{\infty} < \kappa_{\infty}^{-1}$ over $B_S(\boldsymbol{\theta}^*, \min\{A, \frac{\kappa_{\infty}}{M}\})$.

Proof. Define $\alpha_p(\boldsymbol{\theta}) = \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_{SS}^2 L_n(\boldsymbol{\theta}) - \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)]\|_p$ for $p \in \{2, \infty\}$ and $\boldsymbol{\theta} \in B_S(\boldsymbol{\theta}^*, A)$.

Note that for any symmetric matrix \mathbf{A} , we have $\|\mathbf{A}\|_1 = \|\mathbf{A}\|_{\infty}$ and $\|\mathbf{A}\|_2 \leq \sqrt{\|\mathbf{A}\|_1 \|\mathbf{A}\|_{\infty}} \leq \|\mathbf{A}\|_{\infty}$.

Hence by the Assumptions we obtain that when $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \min\{A, \frac{\kappa_{\infty}}{M}\}$ and $p \in \{2, \infty\}$,

$$\alpha_p(\boldsymbol{\theta}) \leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_{\infty} \|\nabla_{SS}^2 L_n(\boldsymbol{\theta}) - \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)\|_{\infty} \leq \frac{1}{2\kappa_{\infty}} M \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \frac{1}{2}.$$

Lemma B.1 leads to

$$\begin{aligned} \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1} - (\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_{\infty} &\leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_{\infty} \frac{\alpha_{\infty}}{1 - \alpha_{\infty}} < \frac{1}{2\kappa_{\infty}}, \\ \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1} - (\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_2 &\leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_2 \frac{\alpha_2}{1 - \alpha_2} < \frac{1}{2\kappa_2}. \end{aligned}$$

Then the proof is finished by triangle's inequality and Assumption B.2. □

Now we are ready to prove Theorem B.1.

Proof of Theorem B.1. First we study the restricted problem $\bar{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathcal{M}} \{L_n(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})\}$.

Take $R(\boldsymbol{\theta}) = \|\boldsymbol{\theta}_{S[p]}\|_1$ and $R^*(\boldsymbol{\theta}) = \|\boldsymbol{\theta}_{S_2}\|_{\infty}$. Let $A_1 = \min\{A, \frac{\kappa_{\infty} \tau}{3M}\}$ and hence $A_1 \leq \min\{A, \frac{\kappa_{\infty}}{M}\}$.

Lemma B.2 shows that $\|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1}\|_2 < \kappa_2^{-1}$ and $\|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}))^{-1}\|_{\infty} < \kappa_{\infty}^{-1}$ over $B_S(\boldsymbol{\theta}^*, A_1)$.

Since $\operatorname{supp}(\boldsymbol{\theta}^*) \subseteq S$, any $\mathbf{h} \in \partial R(\boldsymbol{\theta}^*)$ satisfies $\|\mathbf{h}\|_2 \leq \sqrt{|S_1|}$. Therefore

$$\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda \|\mathbf{h}\|_2 \leq \frac{1}{2} \kappa_2 A_1 \leq \frac{1}{2} \kappa_2 A.$$

Then Corollary A.1 implies that $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{2}{\kappa_2} (\|\nabla_S L(\boldsymbol{\theta}^*)\|_2 + \lambda \sqrt{|S_1|}) \leq A_1$.

Second, we study the L_∞ bound. On the one hand, the optimality condition yields $\nabla_S L_n(\bar{\boldsymbol{\theta}}) \in \lambda \partial \|\bar{\boldsymbol{\theta}}_{[p]}\|_\infty$ and hence $\|\nabla_S L_n(\bar{\boldsymbol{\theta}})\|_\infty \leq \lambda$. On the other hand, by letting $\boldsymbol{\theta}_t = (1-t)\boldsymbol{\theta}^* + t\bar{\boldsymbol{\theta}}$ ($0 \leq t \leq 1$) we have

$$\begin{aligned} \nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*) &= \int_0^1 \nabla_{SS}^2 L_n(\boldsymbol{\theta}_t)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) dt \\ &= \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) + \int_0^1 [\nabla_{SS}^2 L_n(\bar{\boldsymbol{\theta}}_t) - \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) dt. \end{aligned}$$

Hence

$$\begin{aligned} &\|(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) - (\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)]\|_\infty \\ &\leq \int_0^1 \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_{SS}^2 L_n(\bar{\boldsymbol{\theta}}_t) - \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty dt \\ &\leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \sup_{t \in [0,1]} \|\nabla_{SS}^2 L_n(\bar{\boldsymbol{\theta}}_t) - \nabla_{SS}^2 L_n(\boldsymbol{\theta}^*)\|_\infty \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \end{aligned}$$

By Assumptions B.1 and B.2, we obtain that

$$\|(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) - (\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)]\|_\infty \leq \frac{M}{2\kappa_\infty} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.$$

By $\bar{\boldsymbol{\theta}} \in B_S(\boldsymbol{\theta}^*, A_1)$ we have

$$\begin{aligned} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty &\leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{M}{2\kappa_\infty} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \\ &\leq \frac{1}{2\kappa_\infty} (\lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty) + \frac{1}{6} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty. \end{aligned}$$

Therefore,

$$\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq \frac{3}{5\kappa_\infty} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda). \quad (\text{B.3})$$

Third we study the L_1 bound. The bound on $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1$ can be obtained in a similar way. Using the fact that $\|\cdot\|_1 = \|\cdot\|_\infty$ for symmetric matrices,

$$\begin{aligned} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 &\leq \|(\nabla_{SS}^2 L_n(\boldsymbol{\theta}^*))^{-1}\|_1 \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \frac{M}{2\kappa_\infty} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \\ &\leq \frac{1}{2\kappa_\infty} (\lambda |S_1| + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1) + \frac{1}{6} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1. \end{aligned}$$

Hence $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq \frac{3}{5\kappa_\infty} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \lambda |S_1|)$. Since $\text{supp}(\bar{\boldsymbol{\theta}}) \subseteq S$, we also have

$$\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq \sqrt{|S|} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{2\sqrt{|S|}}{\kappa_2} (\|\nabla_S L(\boldsymbol{\theta}^*)\|_2 + \lambda \sqrt{|S_1|}).$$

This gives another L_1 bound.

By Lemma A.1, to derive $\hat{\boldsymbol{\theta}} = \bar{\boldsymbol{\theta}}$ it remains to show that $\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty < \lambda$. Using the Taylor expansion we have

$$\begin{aligned}\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}}) - \nabla_{S_2} L_n(\boldsymbol{\theta}^*) &= \int_0^1 \nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}_t)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) dt \\ &= \nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) + \int_0^1 [\nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}_t) - \nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) dt.\end{aligned}\tag{B.4}$$

On the one hand, the first term in (B.4) follows,

$$\begin{aligned}\|\nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty &= \|[\nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)(\nabla_{S S}^2 L_n(\boldsymbol{\theta}^*))^{-1}][\nabla_{S S}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)]\|_\infty \\ &\leq (1 - \tau) \|\nabla_{S S}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty.\end{aligned}$$

By the Taylor expansion, triangle's inequality, Assumption B.1 and the fact that $\bar{\boldsymbol{\theta}} \in B_S(\boldsymbol{\theta}^*, A_1)$,

$$\begin{aligned}\|\nabla_{S S}^2 L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty &\leq \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \int_0^1 \|[\nabla_{S S}^2 L_n(\bar{\boldsymbol{\theta}}_t) - \nabla_{S S}^2 L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty dt \\ &\leq \|\nabla_S L_n(\bar{\boldsymbol{\theta}})\|_\infty + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + M \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \\ &\leq \lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{\kappa_\infty \tau}{3} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.\end{aligned}$$

On the other hand, we bound the second term in (B.4). Note that $\boldsymbol{\theta}_t \in B_S(\boldsymbol{\theta}^*, A_1)$ for all $t \in [0, 1]$.

Assumption B.1 yields

$$\begin{aligned}\left\| \int_0^1 [\nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}_t) - \nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) dt \right\|_\infty \\ \leq \sup_{t \in [0, 1]} \|\nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}_t) - \nabla_{S_2 S}^2 L_n(\boldsymbol{\theta}^*)\|_\infty \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq \frac{\kappa_\infty \tau}{3} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.\end{aligned}$$

As a result,

$$\begin{aligned}\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty &\leq \|\nabla_{S_2} L_n(\boldsymbol{\theta}^*)\|_\infty + (1 - \tau) \left(\lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{\kappa_\infty \tau}{3} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \right) + \frac{\kappa_\infty \tau}{3} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \\ &\leq \lambda - \tau \left(\lambda - \frac{2\kappa_\infty}{3} \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty - \frac{2}{\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty \right).\end{aligned}$$

Recall that the L_∞ bound in (B.3). By plugging in this estimate, and using the assumptions $0 < \tau < 1$ and $\lambda > \frac{20}{3\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty$, we derive that

$$\begin{aligned}\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty &\leq \lambda - \tau \left(\lambda - \frac{2}{5} (\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda) - \frac{2}{\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty \right) \\ &\leq \lambda - \tau \left(\frac{3}{5} \lambda - \frac{4}{\tau} \|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty \right) < \lambda.\end{aligned}$$

This implies $\hat{\boldsymbol{\theta}} = \bar{\boldsymbol{\theta}}$ and translates all the bounds for $\bar{\boldsymbol{\theta}}$ to the ones for $\hat{\boldsymbol{\theta}}$. The proposition on sign consistency follows from elementary computation, thus we omit its proof. \square

C Proofs of Section 4

C.1 Proof of Theorem 4.1

Proof of Theorem 4.1. Recall that $\hat{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}} \{L_n(\mathbf{y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[p]}\|_1\}$. Also, Assumption 4.3 tells us \mathbf{H}_0 is nonsingular and so is $\mathbf{H} = \begin{pmatrix} \mathbf{I}_p & \mathbf{0}_{p \times K} \\ \mathbf{0}_{K \times p} & \mathbf{H}_0 \end{pmatrix}$. Define $\overline{\mathbf{W}} = \widehat{\mathbf{W}}\mathbf{H}$, $\bar{\boldsymbol{\theta}} = \mathbf{H}^{-1}\hat{\boldsymbol{\theta}}$, $\widehat{\mathbf{B}}_0 = (\mathbf{0}_K^T, \widehat{\mathbf{B}}^T)^T$, $\hat{\boldsymbol{\theta}}^* = \begin{pmatrix} \boldsymbol{\beta}^* \\ \widehat{\mathbf{B}}_0^T \boldsymbol{\beta}^* \end{pmatrix}$ and $\bar{\boldsymbol{\theta}}^* = \mathbf{H}^{-1}\hat{\boldsymbol{\theta}}^*$. We easily see that $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\theta}}_{[p]} = \bar{\boldsymbol{\theta}}_{[p]}$ and $\bar{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}} \{L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[p]}\|_1\}$. Then it follows that $\operatorname{supp}(\hat{\boldsymbol{\beta}}) = \operatorname{supp}(\bar{\boldsymbol{\theta}}_{[p]})$ and $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\| = \|\bar{\boldsymbol{\theta}}_{[p]} - \bar{\boldsymbol{\theta}}_{[p]}^*\| \leq \|\bar{\boldsymbol{\theta}} - \bar{\boldsymbol{\theta}}^*\|$ for any norm $\|\cdot\|$.

Consequently, Theorem 4.1 is reduced to studying $\bar{\boldsymbol{\theta}}$ and the loss function $L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta})$. The Lemma C.1 below implies that all the regularity conditions (with $A = \infty$) in Theorem B.1 are satisfied.

Let w_{tj} and \bar{w}_{tj} be the (i, j) -th element of \mathbf{W} and $\overline{\mathbf{W}}$, respectively. Observe that $L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta}) = \frac{1}{n} \sum_{t=1}^n [-y_t \bar{\mathbf{w}}_t^T \boldsymbol{\theta} + b(\bar{\mathbf{w}}_t^T \boldsymbol{\theta})]$, $\nabla L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta}) = \frac{1}{n} \sum_{t=1}^n [-y_t + b'(\bar{\mathbf{w}}_t^T \boldsymbol{\theta})] \bar{\mathbf{w}}_t$ and $\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^* = \mathbf{X}_1 \boldsymbol{\beta}^*$. Hence $\|\nabla L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty = \varepsilon$ and consequently, $\|\nabla_S L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \leq \varepsilon$, $\|\nabla_S L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_2 \leq \varepsilon \sqrt{|S|}$ and $\|\nabla_S L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_1 \leq \varepsilon |S|$. In addition, $\lambda > 7\varepsilon/\tau \geq \varepsilon$. Based on these estimates, all the results follow from Theorem B.1 and some simple algebra. \square

Here we present the Lemma C.1 used above and its proof.

Lemma C.1. *Let Assumptions 4.1, 4.2 and 4.3 hold. Treat $L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$, and the derivatives below are taken with respect to it. Define $M = M_0^3 M_3 |S|^{3/2}$. Then*

- (i) $\|\nabla_{SS}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\boldsymbol{\theta}) - \nabla_{SS}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \leq M \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2, \forall \boldsymbol{\theta},$
- (ii) $\|(\nabla_{SS}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq \frac{1}{2\kappa_\infty},$
- (iii) $\|(\nabla_{SS}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_2 \leq \frac{1}{2\kappa_2},$
- (iv) $\|\nabla_{S_2S}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq 1 - \tau.$

Proof. (i) Based on the fact that $\mathbf{W}\bar{\boldsymbol{\theta}}^* = \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^* = \mathbf{X}_1 \boldsymbol{\beta}^*$, we have $\nabla^2 L_n(\mathbf{y}, \mathbf{W}\bar{\boldsymbol{\theta}}^*) = \frac{1}{n} \sum_{t=1}^n b''(\bar{\mathbf{w}}_t^T \bar{\boldsymbol{\theta}}^*) \mathbf{w}_t \mathbf{w}_t^T$

and $\nabla^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) = \frac{1}{n} \sum_{t=1}^n b''(\bar{\mathbf{w}}_t^T \bar{\boldsymbol{\theta}}^*) \bar{\mathbf{w}}_t \bar{\mathbf{w}}_t^T$. For any $j, k \in [p+K]$ and $\text{supp}(\boldsymbol{\theta}) \subseteq S$,

$$\begin{aligned} |\nabla_{jk}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\boldsymbol{\theta}) - \nabla_{jk}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)| &\leq \frac{1}{n} \sum_{t=1}^n |b''(\bar{\mathbf{w}}_t^T \boldsymbol{\theta}) - b''(\bar{\mathbf{w}}_t^T \bar{\boldsymbol{\theta}}^*)| \cdot |\bar{\mathbf{w}}_{tj} \bar{\mathbf{w}}_{tk}| \\ &\leq \frac{1}{n} \sum_{t=1}^n M_3 |\bar{\mathbf{w}}_t^T (\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)| \cdot \|\bar{\mathbf{W}}\|_{\max}^2 \end{aligned} \quad (\text{C.1})$$

By the Cauchy-Schwarz inequality and $\|\bar{\mathbf{W}}\|_{\max} \leq \|\mathbf{W}\|_{\max} + \|\bar{\mathbf{W}} - \mathbf{W}\|_{\max} \leq M_0$, we obtain that for $i \in [n]$, $|\bar{\mathbf{w}}_t^T (\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)| = |\bar{\mathbf{w}}_{tS}^T (\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)_S| \leq \|\bar{\mathbf{w}}_{tS}\|_2 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2 \leq \sqrt{|S|} M_0 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2$. Plugging this result back to (C.1), we get

$$\begin{aligned} |\nabla_{jk}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\boldsymbol{\theta}) - \nabla_{jk}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)| &\leq \sqrt{|S|} M_3 M_0^3 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2, \quad \forall j, k \in [p+K]; \\ \|\nabla_{\cdot S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\boldsymbol{\theta}) - \nabla_{\cdot S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_{\infty} &\leq |S|^{3/2} M_3 M_0^3 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2 = M \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2. \end{aligned}$$

(ii) Now we come to the second claim. For any $k \in [p+K]$,

$$\begin{aligned} \|\nabla_{kS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{kS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_{\infty} &\leq \frac{1}{n} \sum_{t=1}^n b''(\mathbf{x}_t^T \boldsymbol{\theta}^*) \|\bar{w}_{tk} \bar{\mathbf{w}}_{tS}^T - w_{tk} \mathbf{w}_{tS}^T\|_{\infty} \\ &\leq \frac{M_2 \sqrt{|S|}}{n} \sum_{t=1}^n \|\bar{w}_{tk} \bar{\mathbf{w}}_{tS}^T - w_{tk} \mathbf{w}_{tS}^T\|_2. \end{aligned}$$

Also, by $\|\mathbf{W}\|_{\max} \leq M_0/2$ and $\|\bar{\mathbf{W}}\|_{\max} \leq M_0$ we have

$$\begin{aligned} \|\bar{w}_{tk} \bar{\mathbf{w}}_{tS}^T - w_{tk} \mathbf{w}_{tS}^T\|_2 &\leq |w_{tk}| \cdot \|(\bar{\mathbf{w}}_{tS} - \mathbf{w}_{tS})^T\|_2 + |\bar{w}_{tk} - w_{tk}| \cdot \|\bar{\mathbf{w}}_{tS}^T\|_2 \\ &\leq \|\mathbf{W}\|_{\max} \|\bar{\mathbf{w}}_{tS} - \mathbf{w}_{tS}\|_2 + |\bar{w}_{tk} - w_{tk}| \cdot \sqrt{|S|} \|\bar{\mathbf{W}}\|_{\max} \\ &\leq \frac{M_0}{2} \|\bar{\mathbf{w}}_{tS} - \mathbf{w}_{tS}\|_2 + M_0 \sqrt{|S|} \cdot |\bar{w}_{tk} - w_{tk}|. \end{aligned}$$

Define $\delta = \max_{j \in [p+K]} (\frac{1}{n} \sum_{t=1}^n |\bar{w}_{tj} - w_{tj}|^2)^{1/2}$. By the Jensen's inequality, $\forall J \subseteq [p+K]$,

$$\frac{1}{n} \sum_{t=1}^n \|\bar{\mathbf{w}}_{tJ} - \mathbf{w}_{tJ}\|_2 \leq \left(\frac{1}{n} \sum_{t=1}^n \|\bar{\mathbf{w}}_{tJ} - \mathbf{w}_{tJ}\|_2^2 \right)^{1/2} \leq \left(\frac{|J|}{n} \max_{j \in [p+K]} \sum_{t=1}^n |\bar{w}_{tj} - w_{tj}|^2 \right)^{1/2} \leq \sqrt{|J|} \delta.$$

As a result,

$$\begin{aligned} \|\nabla_{\cdot S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{\cdot S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_{\infty} \\ = \max_{k \in [p+K]} \|\nabla_{kS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{kS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_{\infty} \leq \frac{3}{2} M_0 M_2 |S| \delta. \end{aligned} \quad (\text{C.2})$$

Let $\alpha = \|(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1} [\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)]\|_{\infty}$. Then

$$\begin{aligned} \alpha &\leq \|(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_{\infty} \|\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_{\infty} \\ &\leq \frac{3}{8\kappa_{\infty}} M_0 M_2 |S| \delta \leq \frac{1}{2}. \end{aligned} \quad (\text{C.3})$$

Lemma B.1 yields

$$\begin{aligned} & \|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \leq \|(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \frac{\alpha}{1-\alpha} \\ & \leq \frac{1}{4\kappa_\infty} \cdot \frac{\alpha}{1-\frac{1}{2}} \leq \frac{3}{16\kappa_\infty^2} M_0 M_2 |S| \delta. \end{aligned}$$

We also have a cruder bound $\|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \leq \frac{1}{4\kappa_\infty}$, which leads to

$$\|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq \|(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty + \frac{1}{4\kappa_\infty} \leq \frac{1}{2\kappa_\infty}. \quad (\text{C.4})$$

(iii) The third argument follows (C.4) easily. Since $\|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_\infty$ holds for any symmetric matrix \mathbf{A} , we have $\|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_2 \leq \frac{1}{4\kappa_\infty} \leq \frac{1}{4\kappa_2}$ and thus $\|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_2 \leq \frac{1}{2\kappa_2}$.

(iv) Finally we prove the last inequality. On the one hand,

$$\begin{aligned} & \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - \nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \\ & \leq \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \\ & \quad + \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)[(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty. \end{aligned}$$

From claim (ii) and (C.2) it is easy to see that

$$\|\nabla_{S_2S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \|(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq \frac{1}{4\kappa_\infty} 3M_0 M_2 |S| \delta.$$

On the other hand, we can take $\mathbf{A} = \nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)$, $\mathbf{B} = \nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)$ and $\mathbf{C} = \nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)$. By Assumption 4.2, $\|\mathbf{A}\mathbf{B}^{-1}\|_\infty \leq 1 - 2\tau \leq 1$. Lemma B.1 forces that

$$\begin{aligned} & \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)[(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty \\ & = \|\mathbf{A}[(\mathbf{B} + \mathbf{C})^{-1} - \mathbf{B}^{-1}]\|_\infty \leq \|\mathbf{A}\mathbf{B}^{-1}\|_\infty \frac{\|\mathbf{C}\mathbf{B}^{-1}\|_\infty}{1 - \|\mathbf{C}\mathbf{B}^{-1}\|_\infty} \leq \frac{\|\mathbf{C}\|_\infty \|\mathbf{B}^{-1}\|_\infty}{1 - \|\mathbf{C}\|_\infty \|\mathbf{B}^{-1}\|_\infty}. \end{aligned}$$

We have shown above in (C.3) that $\|\mathbf{C}\|_\infty \|\mathbf{B}^{-1}\|_\infty \leq \frac{3}{8\kappa_\infty} M_0 M_2 |S| \delta \leq 1/2$. As a result,

$$\|\nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)[(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty \leq \frac{3}{4\kappa_\infty} M_0 M_2 |S| \delta.$$

By combining these estimates, we have

$$\begin{aligned} & \|\nabla_{S_2S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - \nabla_{S_2S}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \\ & \leq \frac{3}{2\kappa_\infty} M_0 M_2 |S| \delta \leq \tau. \end{aligned}$$

Therefore $\|\nabla_{S_2S}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla_{SS}^2 L_n(\mathbf{y}, \bar{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq (1 - 2\tau) + \tau = 1 - \tau$. \square

C.2 Proof of Lemma 4.1

Let $\varepsilon_j = \frac{1}{n} \sum_{t=1}^n \bar{w}_{tj} \eta_t$ for $j \in [p+K]$. Observe that $\varepsilon = \max_{j \in [p+K]} \varepsilon_j$ and $\varepsilon_j = |\frac{1}{n} \sum_{t=1}^n w_{tj} \eta_t + \frac{1}{n} \sum_{t=1}^n (\bar{w}_{tj} - w_{tj}) \eta_t|$. By Cauchy-Schwarz inequality,

$$\left| \frac{1}{n} \sum_{t=1}^n (\bar{w}_{tj} - w_{tj}) \eta_t \right| \leq \left(\frac{1}{n} \sum_{t=1}^n (\bar{w}_{tj} - w_{tj})^2 \right)^{1/2} \left(\frac{1}{n} \sum_{t=1}^n \eta_t^2 \right)^{1/2}.$$

As a result,

$$\varepsilon \leq \max_{j \in [p+K]} \left| \frac{1}{n} \sum_{t=1}^n w_{tj} \eta_t \right| + \left(\frac{1}{n} \sum_{t=1}^n \eta_t^2 \right)^{1/2} \max_{j \in [p+K]} \left(\frac{1}{n} \sum_{t=1}^n (\bar{w}_{tj} - w_{tj})^2 \right)^{1/2}. \quad (\text{C.5})$$

By Theorem 1 and Remark 1 in [Merlevède et al. \(2011\)](#), there exist constants C_1, C_2, C_3 and C_4 such that for any $s \geq 0$,

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{t=1}^n w_{tj} \eta_t \right| > s \right) \leq n \exp \left(-\frac{(ns)^\gamma}{C_1} \right) + \exp \left(-\frac{(ns)^2}{nC_2} \right) + \exp \left(-\frac{(ns)^2}{C_3 n} \exp \left(\frac{(ns)^{\gamma(1-\gamma)}}{C_4 [\log(ns)]^\gamma} \right) \right).$$

From this it is easily seen that for large enough constant $C > 0$, we have $\mathbb{P} \left(\left| \frac{1}{n} \sum_{t=1}^n w_{tj} \eta_t \right| > C \sqrt{\frac{\log p}{n}} \right) \leq p^{-2}$ for all $j \in [p+K]$. Union bounds then force the first in (C.5) to be of order $O_{\mathbb{P}}(\sqrt{\frac{\log p}{n}})$.

Similarly, we can apply the concentration inequality in [Merlevède et al. \(2011\)](#) to get $\frac{1}{n} \sum_{t=1}^n \eta_t^2 = O_{\mathbb{P}}(1)$. It follows from Lemma 3.2 that $\max_{j \in [p+K]} \left[\frac{1}{n} \sum_{t=1}^n (\bar{w}_{tj} - w_{tj})^2 \right]^{1/2} = O_{\mathbb{P}}(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}})$. Hence the second term in (C.5) is of order $O_{\mathbb{P}}(\sqrt{\frac{\log p}{n}} + \frac{1}{\sqrt{p}})$.

C.3 Proof of Lemma 4.2

We follow the proofs of Theorems 2 and 3 in [Fan and Song \(2009\)](#). First of all, the optimality condition for $(\alpha^M, \beta^M, \gamma^M)$ implies that

$$\mathbb{E}\{u[b'(\alpha^M + \mathbf{f}^T \gamma^M + u\beta^M) - y]\} = 0. \quad (\text{C.6})$$

The independence between u and \mathbf{f} as well as the fact that $\mathbb{E}u = 0$ lead to

$$\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \gamma^M)] = 0. \quad (\text{C.7})$$

From (C.7) and $b'' \leq M$ we get

$$|\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \gamma^M + u\beta^M)]| = |\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \gamma^M + u\beta^M)]| - |\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \gamma^M)]|$$

$$\begin{aligned}
&\leq |\mathbb{E}\{u[b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)]\}| \\
&\leq \mathbb{E}[|u| \cdot |b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)|] \\
&\leq \mathbb{E}(|u| \cdot |Mu\beta^M|) = \beta^M M \cdot \mathbb{E} u^2.
\end{aligned}$$

On the other hand, (C.6) leads to $|\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M)]| = |\mathbb{E}(uy)|$. This proves the first part.

When $\mathbb{E}(uy) = 0$, (C.6) yields $\mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M)] = 0$. Then

$$\mathbb{E}\{u[b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)]\} = 0. \quad (\text{C.8})$$

Suppose that $\beta^M \neq 0$. Since b' is strictly increasing, we have

$$\begin{aligned}
0 &< [b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)][(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - (\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)] \\
&= [b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)]u\beta^M
\end{aligned}$$

as long as $u \neq 0$. As $\mathbb{P}(u = 0) = 0$, the inequality above yields

$$0 < \mathbb{E}\{[b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M) - b'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M)]u\beta^M\} = \beta^M \mathbb{E}[ub'(\alpha^M + \mathbf{f}^T \boldsymbol{\gamma}^M + u\beta^M)],$$

which contradicts (C.8).

C.4 Proof of Theorem 4.2

An adaptation of the proof of Theorem 4.1 yields

$$|\widehat{\beta}_j - \beta_j| \leq \|(\widehat{\alpha}_j, \widehat{\boldsymbol{\beta}}_j, \widehat{\gamma}_j^T) - (\alpha_j, \boldsymbol{\beta}_j, \gamma_j^T)\|_2 \leq C' \|\boldsymbol{\varepsilon}_j\|_2 / \kappa, \quad \forall j.$$

Lemma 4.2 show that $|\beta_j| \geq |\mathbb{E}(u_j y)| / (M \cdot \mathbb{E} u_j^2)$. Hence

$$\begin{aligned}
\xi &\leq \rho \min_{j \in \text{supp}(\boldsymbol{\beta}^*)} |\mathbb{E}(u_j y)| / (M \cdot \mathbb{E} u_j^2), \\
\max_{j \in \text{supp}(\boldsymbol{\beta}^*)} \|\boldsymbol{\varepsilon}_j\|_2 &= o_P(\min_{j \in \text{supp}(\boldsymbol{\beta}^*)} |\mathbb{E}(u_j y)| / \mathbb{E} u_j^2)
\end{aligned}$$

lead to $\xi \leq \rho \min_{j \in \text{supp}(\boldsymbol{\beta}^*)} |\beta_j|$ and

$$\max_{j \in \text{supp}(\boldsymbol{\beta}^*)} |\widehat{\beta}_j - \beta_j| \leq C' \max_{j \in \text{supp}(\boldsymbol{\beta}^*)} \|\boldsymbol{\varepsilon}_j\|_2 / \kappa = o_P(\min_{j \in \text{supp}(\boldsymbol{\beta}^*)} |\beta_j|),$$

from where we get $\text{supp}(\boldsymbol{\beta}^*) \subseteq \widehat{S}$.