

Are Latent Factor Regression and Sparse Regression Adequate?

Jianqing Fan

Zhipeng Lou

Mengxin Yu

Abstract

We propose the Factor Augmented sparse linear Regression Model (FARM) that not only encompasses both the latent factor regression and sparse linear regression as special cases but also bridges dimension reduction and sparse regression together. We provide theoretical guarantees for the estimation of our model under the existence of sub-Gaussian and heavy-tailed noises (with bounded $(1 + \vartheta)$ -th moment, for all $\vartheta > 0$) respectively. In addition, the existing works on supervised learning often assume the latent factor regression or sparse linear regression is the true underlying model without justifying its adequacy. To fill in such an important gap, we also leverage our model as the alternative model to test the sufficiency of the latent factor regression and the sparse linear regression models. To accomplish these goals, we propose the **Factor-Adjusted deBiased Test** (FabTest) and a two-stage ANOVA type test respectively. We also conduct large-scale numerical experiments including both synthetic and FRED macroeconomics data to corroborate the theoretical properties of our methods. Numerical results illustrate the robustness and effectiveness of our model against latent factor regression and sparse linear regression models.

Keyword: Factor model, Factor augmented regression, Latent factor regression, Sparse linear regression

¹Jianqing Fan is Frederick L. Moore '18 Professor of Finance, Professor of Statistics, and Professor of Operations Research and Financial Engineering at the Princeton University. Zhipeng Lou is a Postdoctoral Researcher at Department of Operations Research and Financial Engineering, Princeton University. Mengxin Yu is a Ph.D. student at Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA. Emails: {jqfan, zlou, mengxiny}@princeton.edu. The research is supported in part by the NIH grant 2R01-GM072611-16, and the NSF grants DMS-1712591, DMS-2052926, DMS-2053832, and the ONR grant N00014-19-1-2120

1 Introduction

Over the past two decades, along with the development of technology, datasets with high-dimensionality in various fields such as biology, genomics, neuroscience and finance have been collected. One stylized feature of the high-dimensional data is the high co-linearity across features. A common structure to characterize the dependence across features is the approximate factor model [Bai, 2003, Fan et al., 2013], in which the variables are correlated with each other through several common latent factors. More specifically, we assume the observed d -dimensional covariate vector \mathbf{x} follows from the model

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

where \mathbf{f} is a K -dimensional vector of latent factors, $\mathbf{B} \in \mathbb{R}^{d \times K}$ is the corresponding factor loading matrix, and \mathbf{u} is a d -dimensional vector of idiosyncratic component which is uncorrelated with \mathbf{f} .

To tackle the high-dimensionality of datasets, various methods have been proposed. Among these, dimensionality reduction and sparse regression are two popularly used ones to circumvent the curse of dimensionality. They also serve as the backbones for many emerging statistical methods.

In terms of dimension reduction, the factor regression model is one of the most popular methods and has been widely used [Stock and Watson, 2002, Bai and Ng, 2006, Bair et al., 2006, Bai and Ng, 2008, Fan et al., 2017b, Bing et al., 2019, Bunea et al., 2020, Bing et al., 2021]. It assumes that the factors drive both dependent and independent variables as follows:

$$\begin{aligned} Y &= \mathbf{f}^\top \boldsymbol{\gamma} + \varepsilon, \\ \mathbf{x} &= \mathbf{B}\mathbf{f} + \mathbf{u}. \end{aligned} \quad (1.2)$$

Here Y is the response variable and $\varepsilon \in \mathbb{R}$ is the random noise which is independent with the factor \mathbf{f} . When the factors are unobserved, one usually learns the latent factors based on observed \mathbf{x} and substitutes the sample version into the regression model (1.2). There are several methods for estimating latent factors such as Principal Component Analysis (PCA) [Bai, 2003, Fan et al., 2013], maximum likelihood estimation [Bai and Li, 2012], and random projections [Fan and Liao, 2020]. In particular, when the leading Principal Components are used as an estimator for \mathbf{f} , the sample version of (1.2) reduces to the classical Principal Component Regression (PCR) [Hotelling, 1933].

As for sparse regression, a commonly used model is the following (sparse) linear regression:

$$Y = \mathbf{x}^\top \boldsymbol{\beta} + \varepsilon. \quad (1.3)$$

In the high dimensional regime where the dimension d can be much larger than the sample size n , it is commonly assumed that the population parameter vector $\boldsymbol{\beta} \in \mathbb{R}^d$ is sparse. Over the last two decades, various regularized methods, which incorporate this notion of sparsity, have been proposed. See, for instance, LASSO [Tibshirani, 1996], SCAD [Fan and Li, 2001], Least Angle Regression [Efron et al., 2004], Dantzig selector [Candes and Tao, 2007], Adaptive LASSO [Zou, 2006], MCP [Zhang, 2010] and many others. For more details, please refer to Fan et al. [2020b] for a comprehensive survey.

In this paper, we introduce the **Factor Augmented sparse linear Regression Model** (FARM) (1.4), which incorporates both the latent factor and the idiosyncratic component into the covariates,

$$\begin{aligned} Y &= \mathbf{f}^\top \boldsymbol{\gamma}^* + \mathbf{u}^\top \boldsymbol{\beta}^* + \varepsilon, \\ \mathbf{x} &= \mathbf{B}\mathbf{f} + \mathbf{u}, \end{aligned} \quad (1.4)$$

where $\boldsymbol{\gamma}^* \in \mathbb{R}^K$ and $\boldsymbol{\beta}^* \in \mathbb{R}^d$ are population parameter vectors quantifying the contribution of the latent factor \mathbf{f} and the idiosyncratic component \mathbf{u} , respectively. Obviously, the factor regression model (1.2) is a special case of (2.1) in which $\boldsymbol{\beta}^* = 0$. To better illustrate the difference between model (1.4) and the sparse linear model (1.3), our model can be written in an equivalent form,

$$\begin{aligned} Y &= \mathbf{f}^\top \boldsymbol{\varphi}^* + \mathbf{x}^\top \boldsymbol{\beta}^* + \varepsilon, \\ \mathbf{x} &= \mathbf{B}\mathbf{f} + \mathbf{u}, \end{aligned} \quad (1.5)$$

where $\boldsymbol{\varphi}^* = \boldsymbol{\gamma}^* - \mathbf{B}^\top \boldsymbol{\beta}^* \in \mathbb{R}^K$ quantifies the extra contribution of the latent factor \mathbf{f} beyond the observed predictor \mathbf{x} . Therefore, FARM expands the space spanned by \mathbf{x} into useful directions spanned by \mathbf{f} . It is clear that the sparse regression model (1.3) is also a special case of (1.4) with $\boldsymbol{\varphi}^* = 0$. Thus, our model is general enough to bridge the dimensionality reduction and the sparse regression.

The motivation of our factor augmented linear model (1.4) comes from two perspectives.

1. Firstly, it originates from Fan et al. [2020a]. In order to get precise estimation of $\boldsymbol{\beta}^*$ based on highly correlated variables, they study the sparse regression estimation by substituting (1.1) into (1.3)

and obtain

$$Y = (\mathbf{B}\mathbf{f} + \mathbf{u})^\top \boldsymbol{\beta}^* + \varepsilon = \mathbf{f}^\top (\mathbf{B}^\top \boldsymbol{\beta}^*) + \mathbf{u}^\top \boldsymbol{\beta}^* + \varepsilon. \quad (1.6)$$

We observe from (1.6), when the sparse linear regression is adequate, for a given $\boldsymbol{\beta}^*$, the regression coefficient on \mathbf{f} is fixed at $\boldsymbol{\gamma}^* = \mathbf{B}^\top \boldsymbol{\beta}^*$. However, in reality, especially when the variables are highly correlated, it is very likely that the leading factors possess extra contributions to the response instead of only a fixed portion $\mathbf{B}^\top \boldsymbol{\beta}^*$. This results in our proposition of model (1.4), where we augment the leading factors into sparse regression that expands the linear space spanned by \mathbf{x} into useful directions.

2. Secondly, it originates from the factor regression given in (1.2). In reality, the leading common factors \mathbf{f} indeed provides some important contributions to the response, but it is hard to believe that they will have fully explanation power, especially when the effect of the factors is weak. Besides, in real applications, several examples illustrate the poor performance of factor regression model or PCR, see [Jolliffe \[1982\]](#) for more details. Thus, completely ignoring the idiosyncratic component \mathbf{u} will harm in model generalization. This also motivates us to propose model (1.4), in which we augment the sparse regression by incorporating the idiosyncratic component \mathbf{u} into the original factor regression.

In this paper, we first study the properties of estimated parameters under the proposed model (1.4). Specifically, we assume the factors given in (1.4) are unobserved and leverage PCA to estimate them. Incorporated with penalized least-squares with the ℓ_1 -penalty, we derive the ℓ_2 -consistency results for parameter vectors $\boldsymbol{\gamma}^*$ and $\boldsymbol{\beta}^*$. Going beyond the linear regression model and the least squares estimation, our idea can be naturally extended to more general supervised learning models through different loss functions. For instance, quantile regression [[Belloni and Chernozhukov, 2011](#), [Fan et al., 2014](#)], support vector machine [[Zhang et al., 2016](#), [Peng et al., 2016](#)], Huber regression [[Fan et al., 2017a](#), [Sun et al., 2020](#)], generalized linear model [[Van de Geer, 2008](#), [Fan et al., 2020a](#)] and many other variants. In order to demonstrate the general applicability of our proposed methods, in our paper, we further extend our model settings to robust regression. To be more specific, we only assume the existence

of $(1 + \vartheta)$ -th moment of the noise distribution for some $\vartheta > 0$. We adopt Huber loss together with adaptive tuning parameters and ℓ_1 -penalization to derive the consistency results for the parameters of our interest. Besides the aforementioned extensions, it is worth to note that our model is also applicable in the field of causal inference [Imbens and Rubin, 2015, Hernan and Robins, 2019]. To be more specific, the latent factors \mathbf{f} given in our model are able to be treated as the unobserved confounding variables which affect both the covariate \mathbf{x} and the response Y . From the causal perspective, we provide a methodology to conduct (robust) statistical estimation as well as inference of our model under the existence of latent confounding variables.

The aforementioned works on factor regression and sparse linear regression mainly investigate the theoretical properties based on the assumption that either of them is the true underlying model [Stock and Watson, 2002, Tibshirani, 1996, Fan and Li, 2001, Zou, 2006, Bai and Ng, 2006, Zhang, 2010, Fan et al., 2017b, 2020a, Bing et al., 2021]. However, whether a given model is adequate to explain a given dataset plays a crucial role in the model selection step. This motivates us to fill the gap by leveraging our model as the alternative one to perform hypothesis testing on the adequacy of the factor regression model as well as the sparse linear regression model when covariates admit a factor structure.

For the hypothesis test on the adequacy of the latent factor regression model, we consider testing the hypotheses

$$H_0 : Y = \mathbf{f}^\top \boldsymbol{\gamma}^* + \varepsilon \text{ versus } H_1 : Y = \mathbf{f}^\top \boldsymbol{\gamma}^* + \mathbf{u}^\top \boldsymbol{\beta}^* + \varepsilon. \quad (1.7)$$

This amounts to testing $H_0 : \boldsymbol{\beta}^* = 0$ under FARM model. To this end, we propose the **Factor-Adjusted deBiased Test** statistic (FabTest) $\tilde{\boldsymbol{\beta}}_\lambda$ which serves as a de-sparsify version of the estimator $\hat{\boldsymbol{\beta}}_\lambda$ obtained under ℓ_1 -regularization. The asymptotic distribution of the proposed test statistic is derived by leveraging the high-dimensional Gaussian approximation. The critical value controlling the Type-I error is estimated based on the multiplier bootstrap method. As a byproduct, we are also able to conduct entrywise and groupwise hypothesis testing on parameter $\boldsymbol{\beta}^*$ by following similar de-biasing procedure.

For validating the adequacy of the sparse linear regression model, we consider testing the hypothe-

ses

$$H_0 : Y = \mathbf{x}^\top \boldsymbol{\beta}^* + \varepsilon \text{ versus } H_1 : Y = \mathbf{f}^\top \boldsymbol{\varphi}^* + \mathbf{x}^\top \boldsymbol{\beta}^* + \varepsilon, \quad (1.8)$$

or $\boldsymbol{\varphi}^* = 0$ under the FARM model. To tackle the testing problem, we propose a two-stage ANOVA test. In the first stage, we use marginal screening [Fan and Lv, 2008] to pre-select a group of variables which cope well the curse of high dimensionality. In the second stage, we derive the ANOVA-type test statistic. Asymptotic null distribution and the power of the test statistic are derived. In addition, we further extend the aforementioned two-stage ANOVA test to linear multi-modal models [Li and Li, 2021], whose data framework has been well applied in a wide range of scientific fields (e.g multi-omics data in genomics, multimodal neuroimaging data in neuroscience, multimodal electronic health records data in health care).

In summary, our main contributions are as follows:

1. Motivated from the factor regression and sparse regression, we propose the **Factor Augmented (sparse linear) Regression Model** (FARM) (1.4) [also (1.5)] and investigate in the parameter estimation properties on $\boldsymbol{\gamma}^*$ and $\boldsymbol{\beta}^*$ given in (1.4). Our work serves as an extension of Fan et al. [2020a] to a general setting with weaker assumptions. It augments the sparse linear regression in useful directions of common factors.
2. To further demonstrate the wide applicability of our methods, we extend our model to a more robust setting, where we only assume the existence of $(1 + \vartheta)$ -th moment ($\vartheta > 0$) of our noise distribution. Leveraging the ℓ_1 -penalized adaptive Huber estimation, we establish statistical estimation results for our parameters of interest. Comparing with those closely related literature [Fan et al., 2020a, 2021b], our assumption on the moment condition of the noise variable is the weakest. Our robustified factor augmented regression also serves as an extension of Sun et al. [2020] to a more general setting.
3. In terms of testing the adequacy of the factor regression, we propose the **FabTest** by incorporating the factor structure into the de-biased estimators [van de Geer et al., 2014, Zhang and Zhang, 2014, Javanmard and Montanari, 2014]. Accompanied with Gaussian approximation,

the asymptotic distribution of our test statistic is derived. As for implementation, we propose the multiplier bootstrap method to estimate the critical value in order to control the Type-I error.

4. For testing the adequacy of sparse linear regression model, we propose a two stage ANOVA-type testing procedure. Asymptotic distribution (under the null) and power (under the alternative) of our constructed test statistic are investigated. In addition, we further extend the methodology to the multi-modal sparse linear regression model [Li and Li, 2021], by testing whether the sparse linear regression for some given modals is adequate.
5. We conduct large scale simulation studies for our proposed methodology using both synthetic data and real data. Simulation results via synthetic data lend further support to our theoretical findings. As for real data, we apply our methodology to the studies of the macroeconomics dataset named FRED-MD [McCracken and Ng, 2016]. The experimental results also illustrate the high efficiency and robustness of our model (FARM) against latent factor regression as well as sparse linear regression.

1.1 Notation

For a vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_m)^\top \in \mathbb{R}^m$, we denote its ℓ_q norm as $\|\boldsymbol{\gamma}\|_q = (\sum_{\ell=1}^m |\gamma_\ell|^q)^{1/q}$, $1 \leq q < \infty$, and write $\|\boldsymbol{\gamma}\|_\infty = \max_{1 \leq \ell \leq m} |\gamma_\ell|$. For any integer m , we denote $[m] = \{1, \dots, m\}$. The sub-Gaussian norm of a scalar random variable Z is defined as $\|Z\|_{\psi_2} = \inf\{t > 0 : \mathbb{E} \exp(Z^2/t^2) \leq 2\}$. For a random vector $\mathbf{x} \in \mathbb{R}^m$, we use $\|\mathbf{x}\|_{\psi_2} = \sup_{\|\mathbf{v}\|_2=1} \|\mathbf{v}^\top \mathbf{x}\|_{\psi_2}$ to denote its sub-Gaussian norm. Let $\mathbb{I}\{\cdot\}$ denote the indicator function and let \mathbf{I}_K denotes the identity matrix in $\mathbb{R}^{K \times K}$. For a matrix $\mathbf{A} = [A_{jk}]$, we define $\|\mathbf{A}\|_{\mathbb{F}} = \sqrt{\sum_{jk} A_{jk}^2}$, $\|\mathbf{A}\|_{\max} = \max_{jk} |A_{jk}|$ and $\|\mathbf{A}\|_\infty = \max_j \sum_k |A_{jk}|$ to be its Frobenius norm, element-wise max-norm and matrix ℓ_∞ -norm, respectively. Moreover, we use $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ to denote the minimal and maximal eigenvalues of \mathbf{A} , respectively. We use $|\mathcal{A}|$ to denote the cardinality of set \mathcal{A} . For two positive sequences $\{a_n\}_{n \geq 1}$, $\{b_n\}_{n \geq 1}$, we write $a_n = O(b_n)$ if there exists a positive constant C such that $a_n \leq C \cdot b_n$ and we write $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$. In addition, $a_n = O_{\mathbb{P}}(b_n)$ and $a_n = o_{\mathbb{P}}(b_n)$ have similar meanings as above except that the relationship of a_n/b_n holds with high probability.

1.2 RoadMap

The rest of this paper is organized as follows. We study the parameter estimation properties of our proposed model (FARM) in section 2, where theoretical results of both regular and robust estimators are analyzed. In section 3, we construct a de-biased test statistic to test the adequacy of latent factor regression model. In addition, in section 4, we construct a two-stage ANOVA test to study the adequacy of sparse linear regression under the setting with highly correlated features. Moreover, to corroborate our theoretical findings, in section 5, we conduct exhaustive simulation studies. Last but not least, we apply our methodology to study the real data FRED-MD in section 5.4.

2 Factor Augmented Regression Model

The primary objective of this section is to propose a regularized estimation method for our factor augmented sparse linear model and investigate the corresponding statistical properties. Suppose we observe n independent and identically distributed (i.i.d.) random samples $\{(\mathbf{x}_t, Y_t)\}_{t=1}^n$ from (\mathbf{x}, Y) , which satisfy that

$$\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t \text{ and } Y_t = \mathbf{f}_t^\top \boldsymbol{\gamma}^* + \mathbf{u}_t^\top \boldsymbol{\beta}^* + \varepsilon_t, \quad t = 1, \dots, n, \quad (2.1)$$

where $\mathbf{f}_1, \dots, \mathbf{f}_n \in \mathbb{R}^K$, $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{R}^d$ and $\varepsilon_1, \dots, \varepsilon_n \in \mathbb{R}$ are i.i.d. realizations of \mathbf{f} , \mathbf{u} and ε , respectively. To ease the presentation, we rewrite (2.1) in a more compact matrix form as follows,

$$\begin{aligned} \mathbf{X} &= \mathbf{F}\mathbf{B}^\top + \mathbf{U}, \\ \mathbf{Y} &= \mathbf{F}\boldsymbol{\gamma}^* + \mathbf{U}\boldsymbol{\beta}^* + \mathcal{E}, \end{aligned} \quad (2.2)$$

where $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$, $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_n)^\top$, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_n)^\top$, $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\mathcal{E} = (\varepsilon_1, \dots, \varepsilon_n)^\top$. Throughout the whole paper, we assume we only get access to observations $\{(\mathbf{x}_t, Y_t)\}_{t=1}^n$. Both the latent factors \mathbf{F} and the idiosyncratic components \mathbf{U} are unobserved and need to be estimated from the observed predictors \mathbf{X} . Thus, in the following, we shall first illustrate how to estimate \mathbf{F} and \mathbf{U} and then proceed with the regularized estimation for model (2.2).

2.1 Factor Estimation

Since only the predictor vector \mathbf{x} is observable, the latent factor \mathbf{f} and the corresponding loading matrix \mathbf{B} are not identifiable under the factor model (1.1). More specifically, for any non-singular matrix $\mathbf{S} \in \mathbb{R}^{K \times K}$, we have $\mathbf{x} = \mathbf{B}\mathbf{f} + \mathbf{u} = (\mathbf{B}\mathbf{S})(\mathbf{S}^{-1}\mathbf{f}) + \mathbf{u}$. To resolve this issue, we impose the following identifiability conditions [Bai, 2003, Fan et al., 2013]:

$$\text{Cov}(\mathbf{f}) = \mathbf{I}_K \text{ and } \mathbf{B}^\top \mathbf{B} \text{ is diagonal.}$$

Consequently, the constrained least squares estimator of (\mathbf{F}, \mathbf{B}) based on \mathbf{X} is given by

$$\begin{aligned} (\hat{\mathbf{F}}, \hat{\mathbf{B}}) &= \arg \min_{\mathbf{F} \in \mathbb{R}^{n \times K}, \mathbf{B} \in \mathbb{R}^{d \times K}} \|\mathbf{X} - \mathbf{F}\mathbf{B}^\top\|_{\mathbb{F}}^2 \\ \text{subject to } &\frac{1}{n}\mathbf{F}^\top \mathbf{F} = \mathbf{I}_K \text{ and } \mathbf{B}^\top \mathbf{B} \text{ is diagonal.} \end{aligned}$$

Elementary manipulation yields that the columns of $\hat{\mathbf{F}}/\sqrt{n}$ are the eigenvectors corresponding to the largest K eigenvalues of the matrix $\mathbf{X}\mathbf{X}^\top$ and $\hat{\mathbf{B}} = (\hat{\mathbf{F}}^\top \hat{\mathbf{F}})^{-1} \hat{\mathbf{F}}^\top \mathbf{X} = n^{-1} \hat{\mathbf{F}}^\top \mathbf{X}$. Then the least squares estimator for \mathbf{U} is given by $\hat{\mathbf{U}} = \mathbf{X} - \hat{\mathbf{F}}\hat{\mathbf{B}}^\top = (\mathbf{I}_n - n^{-1} \hat{\mathbf{F}}\hat{\mathbf{F}}^\top) \mathbf{X}$.

Before presenting the asymptotic properties of the estimators $\{\hat{\mathbf{F}}, \hat{\mathbf{B}}, \hat{\mathbf{U}}\}$, we first impose some regularity conditions.

Assumption 2.1. There exists a positive constant $c_0 < \infty$ such that $\|\mathbf{f}\|_{\psi_2} \leq c_0$ and $\|\mathbf{u}\|_{\psi_2} \leq c_0$.

Assumption 2.2. There exists a constant $\tau > 1$ such that $d/\tau \leq \lambda_{\min}(\mathbf{B}^\top \mathbf{B}) \leq \lambda_{\max}(\mathbf{B}^\top \mathbf{B}) \leq d\tau$.

Assumption 2.3. Let $\Sigma = \text{Cov}(\mathbf{u})$. There exists a constant $\Upsilon > 0$ such that $\|\mathbf{B}\|_{\max} \leq \Upsilon$ and

$$\mathbb{E}|\mathbf{u}^\top \mathbf{u} - \text{tr}(\Sigma)|^4 \leq \Upsilon d^2.$$

Assumption 2.4. There exist a positive constant $\kappa < 1$ such that $\kappa \leq \lambda_{\min}(\Sigma)$, $\|\Sigma\|_1 \leq 1/\kappa$ and $\min_{1 \leq k, \ell \leq d} \text{Var}(u_k u_\ell) \geq \kappa$.

Remark 1. Assumptions 2.1–2.4 are standard assumptions in the studies of large dimensional factor model. We refer to Bai [2003], Fan et al. [2013] and Li et al. [2018] for more details. \square

We next summarize the theoretical results related to consistent factor estimation in the following proposition which directly follows from Lemmas D.1 and D.2 in [Wang and Fan \[2017\]](#).

Proposition 2.1. *Assume that $\log n = o(d)$. Let $\mathbf{H} = n^{-1}\mathbf{V}^{-1}\widehat{\mathbf{F}}^\top \mathbf{F}\mathbf{B}^\top \mathbf{B}$, where $\mathbf{V} \in \mathbb{R}^{K \times K}$ is a diagonal matrix consisting of the first K largest eigenvalues of the matrix $n^{-1}\mathbf{X}\mathbf{X}^\top$. Then, under Assumptions 2.1–2.4, we have*

1. $\|\widehat{\mathbf{F}} - \mathbf{F}\mathbf{H}^\top\|_{\mathbb{F}}^2 = O_{\mathbb{P}}(n/d + 1/n)$.
2. For any $\mathcal{I} \subset \{1, 2, \dots, d\}$, we have $\max_{\ell \in \mathcal{I}} \sum_{t=1}^n |\widehat{u}_{t\ell} - u_{t\ell}|^2 = O_{\mathbb{P}}(\log |\mathcal{I}| + n/d)$.
3. $\|\mathbf{H}^\top \mathbf{H} - \mathbf{I}_K\|_{\mathbb{F}}^2 = O_{\mathbb{P}}(1/n + 1/d)$.
4. $\max_{\ell \in [d]} \|\widehat{\mathbf{b}}_\ell - \mathbf{H}\mathbf{b}_\ell\|_2^2 = O_{\mathbb{P}}\{(\log d)/n\}$.

Remark 2. In practice, the number of latent factors K is typically unknown and it is an important issue to determine K in a data-driven way. There have been various methods proposed in the literature to estimate the number K [[Bai and Ng, 2002](#), [Lam and Yao, 2012](#), [Ahn and Horenstein, 2013](#), [Fan et al., 2021a](#)]. Our theories always work as long as we replace K by any consistent estimator \widehat{K} , i.e. we only require

$$\mathbb{P}(\widehat{K} = K) \rightarrow 1, \text{ as } n \rightarrow \infty.$$

Thus, without loss of generality, we assume the number of factors K is known throughout all the theories developed in this paper. As for the application part, throughout this paper, we utilize the eigenvalue ratio method [[Lam and Yao, 2012](#), [Ahn and Horenstein, 2013](#)] to select the number of factors. More specifically, we let $\lambda_k(\mathbf{X}\mathbf{X}^\top)$ denote the eigenvalues of the Gram matrix $\mathbf{X}\mathbf{X}^\top$ and the number of factors is given by

$$\widehat{K} = \arg \max_{K \leq \mathcal{K}} \frac{\lambda_k(\mathbf{X}\mathbf{X}^\top)}{\lambda_{k+1}(\mathbf{X}\mathbf{X}^\top)},$$

where $1 \leq \mathcal{K} \leq n$ is a prescribed upper bound for K . □

2.2 Regularization Estimation

Under the high dimensional regime where the dimension d can be much larger than the sample size n , it is often assumed that only a small portion of the predictors contribute to the response variable, which amounts to assuming that the true parameter vector β^* is sparse. Then the regularized estimator for the unknown parameter vectors β^* and γ^* of our factor augmented linear model is defined as follows:

$$(\hat{\beta}_\lambda, \hat{\gamma}) = \arg \min_{\beta \in \mathbb{R}^d, \gamma \in \mathbb{R}^K} \left\{ \frac{1}{2n} \|\mathbf{Y} - \hat{\mathbf{U}}\beta - \hat{\mathbf{F}}\gamma\|_2^2 + \lambda \|\beta\|_1 \right\}, \quad (2.3)$$

where $\lambda > 0$ is a tuning parameter.

We let $\tilde{\mathbf{Y}} = (\mathbf{I}_n - \hat{\mathbf{P}})\mathbf{Y}$ denote the residuals of the response vector \mathbf{Y} after projecting onto the column space of $\hat{\mathbf{F}}$, where $\hat{\mathbf{P}} = n^{-1}\hat{\mathbf{F}}\hat{\mathbf{F}}^\top$ is the corresponding projection matrix. Recall that $\hat{\mathbf{U}} = (\mathbf{I}_n - \hat{\mathbf{P}})\mathbf{X}$. Hence $\hat{\mathbf{F}}^\top \hat{\mathbf{U}} = 0$ and it is straightforward to verify that the solution of (2.3) is equivalent to

$$\begin{aligned} \hat{\beta}_\lambda &= \arg \min_{\beta \in \mathbb{R}^d} \left\{ \frac{1}{2n} \|\tilde{\mathbf{Y}} - \hat{\mathbf{U}}\beta\|_2^2 + \lambda \|\beta\|_1 \right\}, \\ \hat{\gamma} &= (\hat{\mathbf{F}}^\top \hat{\mathbf{F}})^{-1} \hat{\mathbf{F}}^\top \mathbf{Y} = \frac{1}{n} \hat{\mathbf{F}}^\top \mathbf{Y}. \end{aligned}$$

For any subset \mathcal{S} of $\{1, \dots, d\}$, we define the convex cone $\mathcal{C}(\mathcal{S}, 3) = \{\delta \in \mathbb{R}^d : \|\delta_{\mathcal{S}^c}\|_1 \leq 3\|\delta_{\mathcal{S}}\|_1\}$. For simplicity of notation, we write

$$\mathcal{V}_{n,d} = \frac{n}{d} + \sqrt{\frac{\log d}{n}} + \sqrt{\frac{n \log d}{d}}. \quad (2.4)$$

To investigate the consistency property of $(\hat{\beta}_\lambda, \hat{\gamma})$, we impose the following moment condition on the random noise ε .

Assumption 2.5. There exists a positive constant $c_1 < \infty$ such that $\|\varepsilon\|_{\psi_2} \leq c_1$.

Theorem 2.2. Recall $\varphi^* = \gamma^* - \mathbf{B}^\top \beta^* \in \mathbb{R}^K$. Under Assumptions 2.1–2.5, we have

$$\|\hat{\gamma} - \mathbf{H}\gamma^*\|_2 = O_{\mathbb{P}} \left\{ \frac{1}{\sqrt{n}} + \left(\frac{1}{\sqrt{n}} + \frac{1}{\sqrt{d}} \right) \|\varphi^*\|_2 + \|\beta^*\|_1 \left(\sqrt{\frac{\log |\mathcal{S}_*|}{n}} + \frac{1}{\sqrt{d}} \right) \right\},$$

where $\mathcal{S}_* = \{j \in [d] : \beta_j^* \neq 0\}$ and $|\mathcal{S}_*|$ is its cardinality. Furthermore, if $|\mathcal{S}_*| \left(\frac{\log d}{n} + \frac{1}{d} \right) \rightarrow 0$, then, by taking $\lambda = (\mathcal{I}_0/n) \|\widehat{\mathbf{U}}^\top (\widetilde{\mathbf{Y}} - \widehat{\mathbf{U}}\beta^*)\|_\infty$ for some constant $\mathcal{I}_0 \geq 2$, we have $\widehat{\beta}_\lambda - \beta^* \in \mathcal{C}(\mathcal{S}_*, 3)$ and

$$\|\widehat{\beta}_\lambda - \beta^*\|_2 = O_{\mathbb{P}} \left(\sqrt{\frac{|\mathcal{S}_*| \log p}{n}} + \frac{\mathcal{V}_{n,d} \|\varphi^*\|_2 \sqrt{|\mathcal{S}_*|}}{n} \right). \quad (2.5)$$

Remark 3. In most of literature investigating the regularized estimation of sparse linear regression model (1.3), it is commonly assumed that the observed covariate vector \mathbf{x} is a sub-Gaussian random vector with bounded sub-Gaussian norm $\|\mathbf{x}\|_{\psi_2}$. See, for instance, [Loh and Wainwright \[2012\]](#), [Nickl and Van De Geer \[2013\]](#), [van de Geer et al. \[2014\]](#), [Zhang and Cheng \[2017\]](#) and many others. However, such assumption can be unreasonable in the presence of highly correlated covariates. To see this, suppose now both \mathbf{f} and \mathbf{u} are Gaussian random vectors and the underlying \mathbf{x} satisfies the factor model (1.1). Then \mathbf{x} is also a Gaussian random vector with $\text{Cov}(\mathbf{x}) = \mathbf{B}\mathbf{B}^\top + \Sigma$. Under the pervasiveness condition (Assumption 2.2) and Assumption 2.4, it is straightforward to verify that $\|\mathbf{x}\|_{\psi_2} = \sqrt{8/3} \lambda_{\max}(\mathbf{B}\mathbf{B}^\top + \Sigma) \asymp d$, which violates the assumption on bounded sub-Gaussian norm. In contrast, our model can circumvent such issue because we decompose the covariate \mathbf{x} into (\mathbf{f}, \mathbf{u}) , and we only need impose sub-Gaussian assumption on (\mathbf{f}, \mathbf{u}) . As the sparse linear regression model serves as a special case to our model, our model serves as a more robust choice to conduct parameter estimation comparing with using linear regression directly, even if the sparse linear regression model is adequate. \square

Remark 4. Theorem 2.2 substantially generalize the results in [Fan et al. \[2020a\]](#) with weaker assumptions. First, we did not impose the irrepresentable condition on the design matrix \mathbf{U} , only the lower bound on $\Sigma = \text{Cov}(\mathbf{u})$ is required. In addition, although [Fan et al. \[2020a\]](#) also decompose the covariate \mathbf{x} into (\mathbf{f}, \mathbf{u}) in order to get precise estimator for β^* , they mainly focus on the linear model $Y = \mathbf{x}^\top \beta^* + \varepsilon$ which corresponds to the special case with $\varphi^* = 0$ in our results given in Theorem 2.2. \square

Remark 5. We note that our study is very different from the related work by [Fan et al. \[2021b\]](#), although they also study one kind of factor augment linear regression model. To be more specific, they assume the response $Y_{i,t}$ is given in a penal form with $i \in [N], t \in [T]$, which is generated from the

model $Y_{i,t} = \boldsymbol{\gamma}_i^\top \mathbf{f}_t + u_{i,t}$. Here $u_{i,t}, i \in [N], t \in [T]$ is the idiosyncratic component. They incorporate the sparse linear regression into their model by assuming $u_{i,t} = \boldsymbol{\beta}_i^\top \mathbf{u}_{-i,t} + \epsilon_{i,t}, \forall i \in [N]$. Thus, their factor augmented sparse linear regression model heavily relies on the penal data structure. In contrast, we study the cross-sectional data and focus on different inference problems.

□

2.3 Factor Augmented Robust Linear Regression

In reality, datasets, especially collected from the field of finance, are often contaminated by noises with relatively heavy tails. To resolve such issue, we leverage the adaptive Huber regression to study the parameter of interest in our FARM under the existence of heavy-tailed noise [Sun et al., 2020].

We first introduce some notation and basic definitions. Let $\rho_\omega(\cdot)$ denote the Huber function,

$$\rho_\omega(z) = \begin{cases} z^2/2, & \text{if } |z| \leq \omega, \\ \omega z - \omega^2/2, & \text{if } |z| > \omega, \end{cases}$$

where $\omega > 0$ is the robustification parameter which balances robustness and bias. Following the intuition of (2.3), our factor augmented Huber estimator for $(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*)$ is given by

$$(\hat{\boldsymbol{\beta}}_h, \hat{\boldsymbol{\gamma}}_h) = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^d, \boldsymbol{\gamma} \in \mathbb{R}^K} \left\{ \frac{1}{n} \sum_{t=1}^n \rho_\omega(y_t - \hat{\mathbf{u}}_t^\top \boldsymbol{\beta} - \hat{\mathbf{f}}_t^\top \boldsymbol{\gamma}) + \lambda \|\boldsymbol{\beta}\|_1 \right\}, \quad (2.6)$$

where $\lambda > 0$ is a tuning parameter. For simplicity of notation, we write $\hat{\boldsymbol{\phi}}_h = (\hat{\boldsymbol{\beta}}_h^\top, \hat{\boldsymbol{\gamma}}_h^\top)^\top \in \mathbb{R}^{d+K}$ and $\tilde{\boldsymbol{\phi}} = (\boldsymbol{\beta}^{*\top}, \tilde{\boldsymbol{\gamma}}^\top)^\top \in \mathbb{R}^{d+K}$, where $\tilde{\boldsymbol{\gamma}} = \hat{\mathbf{B}}^\top \boldsymbol{\beta}^* + n^{-1} \hat{\mathbf{F}}^\top \mathbf{F} \boldsymbol{\varphi}^*$. The following theorem establishes the statistical consistency of $\hat{\boldsymbol{\phi}}_h$.

Proposition 2.3. *Assume that $\mathbb{E}|\varepsilon|^{1+\vartheta} < \infty$ for some constant $\vartheta > 0$. Let*

$$\omega \asymp \left(\frac{n}{\log d} \right)^{\frac{1}{1+(\vartheta+1)}} \quad \text{and} \quad \lambda \asymp \left(\frac{\log d}{n} \right)^{\frac{\vartheta \wedge 1}{1+(\vartheta+1)}}.$$

Furthermore, we assume that $(|\mathcal{S}_*| + K)(\log d)^{3/2} = o(n)$,

$$\frac{\log n}{n + \sqrt{d}} \|\boldsymbol{\varphi}^*\|_2 = o(\omega) \quad \text{and} \quad \mathcal{V}_{n,d} \|\boldsymbol{\varphi}^*\|_2 = O(\omega \log d). \quad (2.7)$$

Then, under Assumptions 2.1–2.4, we have

$$\|\hat{\boldsymbol{\phi}}_h - \tilde{\boldsymbol{\phi}}\|_1 = O_{\mathbb{P}} \left\{ (|\mathcal{S}_*| + K) \left(\frac{\log d}{n} \right)^{\frac{\vartheta \wedge 1}{1 + (\vartheta \wedge 1)}} \right\}.$$

We establish the ℓ_1 -statistical rate for our parameters in model (1.4)[also (1.5)] by only assuming the existence of $(1 + \vartheta)$ -th moment of the noise distribution. Specifically, when $\vartheta \geq 1$, the results reduce to the same rates as the sub-Gaussian assumption of ε . Our result serves as an extension of [Sun et al. \[2020\]](#) to a more general setting by incorporating latent factors.

3 Is Factor Regression Model Adequate?

The latent factor regression is widely applied in many fields as an efficient dimension reduction method. A natural question arises is whether the model is adequate and FARM (1.4) serves naturally as the alternative model. To be more specific, we consider testing the hypotheses

$$H_0 : \boldsymbol{\beta}^* = 0 \text{ versus } H_1 : \boldsymbol{\beta}^* \neq 0 \quad (3.1)$$

in FARM (1.4). As the penalized least-squares estimator $\hat{\boldsymbol{\beta}}_\lambda$ is used for estimating $\boldsymbol{\beta}^*$, it creates biases and make it difficult for inferences. Thus, we first introduce a de-biased version of $\hat{\boldsymbol{\beta}}_\lambda$ given in (2.3).

3.1 Bias Correction

We begin with the construction of bias-corrected estimator for $\boldsymbol{\beta}^*$ following similar idea of [Zhang and Zhang \[2014\]](#), [van de Geer et al. \[2014\]](#) and [Javanmard and Montanari \[2014\]](#). Specifically, let $\hat{\boldsymbol{\Theta}} \in \mathbb{R}^{d \times d}$ be an approximation for the inverse of the Gram matrix $\tilde{\boldsymbol{\Sigma}} = n^{-1} \hat{\mathbf{U}}^\top \hat{\mathbf{U}}$, the de-biased estimator for $\boldsymbol{\beta}^*$ is then defined as

$$\tilde{\boldsymbol{\beta}}_\lambda = \hat{\boldsymbol{\beta}}_\lambda + \frac{1}{n} \hat{\boldsymbol{\Theta}} \hat{\mathbf{U}}^\top (\mathbf{Y} - \hat{\mathbf{U}} \hat{\boldsymbol{\beta}}_\lambda). \quad (3.2)$$

The rationale behind such construction is that we are able to decompose estimation error as

$$\tilde{\boldsymbol{\beta}}_\lambda - \boldsymbol{\beta}^* = \frac{1}{n} \hat{\boldsymbol{\Theta}} \hat{\mathbf{U}}^\top \mathcal{E} + \frac{1}{n} \hat{\boldsymbol{\Theta}} \hat{\mathbf{U}}^\top \mathbf{F} \boldsymbol{\varphi}^* + (\mathbf{I}_d - \hat{\boldsymbol{\Theta}} \tilde{\boldsymbol{\Sigma}}) (\hat{\boldsymbol{\beta}}_\lambda - \boldsymbol{\beta}^*), \quad (3.3)$$

after we expand \mathbf{Y} according to (2.2) and replace \mathbf{X} by $\mathbf{X} = \widehat{\mathbf{F}}\widehat{\mathbf{B}} + \widehat{\mathbf{U}}$. The first term on the right hand side of (3.3) quantifies the uncertainty of our estimator $\tilde{\beta}_\lambda$ and the last two terms are biases which will be shown to be of smaller order.

One observes that constructing the de-biased estimator $\tilde{\beta}_\lambda$ given above requires an estimator $\widehat{\Theta}$. There exist many methodologies devoting to estimating such precision matrix, for example, the node-wise regression proposed in [Zhang and Zhang \[2014\]](#) and [van de Geer et al. \[2014\]](#), and the CLIME-type estimator given in [Cai et al. \[2011\]](#), [Javanmard and Montanari \[2014\]](#) and [Avella-Medina et al. \[2018\]](#). In our work, we do not restrict $\widehat{\Theta}$ to be any specific one, but require to satisfy the following general conditions.

Assumption 3.1. Let $\Theta = \Sigma^{-1}$ with Σ defined in Assumption 2.3. There exist positive Λ_{\max} and Δ_∞ such that

$$\|\mathbf{I}_d - \widehat{\Theta}\tilde{\Sigma}\|_{\max} = O_{\mathbb{P}}(\Lambda_{\max}) \text{ and } \|\widehat{\Theta} - \Theta\|_\infty = O_{\mathbb{P}}(\Delta_\infty).$$

Without loss of generality, here we assume that $\Delta_\infty \leq \|\Theta\|_\infty$.

Remark 6. To give a concrete example, under the mild conditions therein, Assumption 3.1 is satisfied with

$$\Lambda_{\max} = O\left(\sqrt{\frac{\log d}{n}} + \frac{1}{\sqrt{d}}\right) \text{ and } \Delta_\infty = O\left(\max_{j \in [d]} |\mathcal{S}_j| \sqrt{\frac{\log d}{n} + \frac{1}{d}}\right),$$

by using node-wise regression [[Zhang and Zhang, 2014](#), [van de Geer et al., 2014](#)], where $|\mathcal{S}_j| = \sum_{k=1}^d \mathbb{I}\{\Theta_{jk} \neq 0\}$ quantifies the sparsity of j -th column of the precision matrix Θ for each $1 \leq j \leq d$. In Appendix C.1, we will provide a detailed analysis on estimating $\tilde{\Sigma}^{-1}$ via node-wise regression and establish precise theoretical upper bounds for the statistical rates given in Assumption 3.1. \square

3.2 Gaussian Approximation

The goal of this section is to derive the asymptotic distribution of $\|\tilde{\beta}_\lambda - \beta^*\|_\infty$ in the high dimensional setting. To this end, we apply the Gaussian approximation result given in [Chernozhukov et al. \[2013, 2017, 2020\]](#) for high dimensional random vectors. More specifically, we let $\mathbf{Z} = (Z_1, \dots, Z_d)^\top \in \mathbb{R}^d$

be a zero-mean Gaussian random vector with the same covariance matrix as that of $n^{-1/2}\Theta\mathbf{U}^\top\mathcal{E}$, that is,

$$\text{Cov}(\mathbf{Z}) = \text{Cov}\left(\frac{1}{\sqrt{n}}\Theta\mathbf{U}^\top\mathcal{E}\right) = \sigma^2\Theta. \quad (3.4)$$

We next present the theoretical results on Gaussian approximation of our test statistics under some mild conditions.

Theorem 3.1. *Recall $\boldsymbol{\varphi}^* = \boldsymbol{\gamma}^* - \mathbf{B}^\top\boldsymbol{\beta}^* \in \mathbb{R}^K$. We assume that $(\log d)^5/n \rightarrow 0$,*

$$(\Lambda_{\max}|\mathcal{S}_\star| + \Delta_\infty)\log d \rightarrow 0 \text{ and } \left(\mathcal{V}_{n,d}\|\boldsymbol{\varphi}^*\|_2 + \sqrt{\frac{n}{d}} + \sqrt{\log d}\right)\|\Theta\|_\infty\sqrt{\frac{\log d}{n}} \rightarrow 0, \quad (3.5)$$

with $\mathcal{V}_{n,d}$ given by (2.4). Then under Assumption 3.1, we have

$$\sup_{x>0} \left| \mathbb{P}\left(\sqrt{n}\|\tilde{\boldsymbol{\beta}}_\lambda - \boldsymbol{\beta}^*\|_\infty \leq x\right) - \mathbb{P}(\|\mathbf{Z}\|_\infty \leq x) \right| \rightarrow 0.$$

For any $\alpha \in (0, 1)$, let $c_{1-\alpha}$ denote the $(1-\alpha)$ -th quantile of the distribution of $\|\mathbf{Z}\|_\infty$. Theorem 3.1 leads to an approximately level α test for (3.1) as follows:

$$\psi_{\infty,\alpha} = \mathbb{I}\left\{\sqrt{n}\|\tilde{\boldsymbol{\beta}}_\lambda\|_\infty > c_{1-\alpha}\right\}. \quad (3.6)$$

3.3 Gaussian multiplier bootstrap

The critical value $c_{1-\alpha}$ depends on the unknown σ^2 and Θ , which can be estimated by the following Gaussian multiplier bootstrap.

1. Generate i.i.d. random variables $\xi_1, \dots, \xi_n \sim N(0, 1)$ and compute

$$\hat{L} = \frac{1}{\sqrt{n}}\|\hat{\Theta}\hat{\mathbf{U}}^\top\xi\|_\infty, \text{ where } \boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_n)^\top.$$

2. Repeat the first step independently for B times and obtain $\hat{L}_1, \dots, \hat{L}_B$. Estimate the critical value $c_{1-\alpha}$ via $1 - \alpha$ quantile of the empirical distribution of the bootstrap statistics:

$$\hat{c}_{1-\alpha} = \inf\{t \geq 0 : H_B(t) \geq 1 - \alpha\}, \text{ where } H_B(t) = \frac{1}{B} \sum_{b=1}^B \mathbb{I}\left\{\hat{L}_b \leq t\right\}.$$

Reject the null hypothesis H_0 when $\sqrt{n}\|\tilde{\beta}_\lambda\|_\infty/\hat{\sigma} > \hat{c}_{1-\alpha}$, for a given consistent estimator $\hat{\sigma}$ of σ . To validate the procedure, we need some additional conditions on $\hat{\Theta}$ and $\hat{\sigma}$.

Assumption 3.2. There exists a $\Delta_{\max} > 0$ such that $\|\hat{\Theta} - \Theta\|_{\max} = O_{\mathbb{P}}(\Delta_{\max})$.

Assumption 3.3. There exists a $0 < \Delta_\sigma \leq 1$ such that $|\hat{\sigma}/\sigma - 1| = O_{\mathbb{P}}(\Delta_\sigma)$.

Remark 7. The estimation of σ^2 for high dimensional linear regression has been extensively in the literature. For example, [Fan et al. \[2012\]](#) proposed refitted cross-validation to construct a consistent estimator with clearly quantified uncertainty of $\hat{\sigma}$ in ultra-high dimension. In addition, [Sun and Zhang \[2012\]](#) and [Yu and Bien \[2019\]](#) derived scaled-Lasso and organic Lasso respectively for estimating σ . Like our case of estimating Θ , we also do not restrict estimating σ by any fixed method mentioned above, our theory works as long as the general condition of Assumption 3.3 holds. \square

Let $\mathbb{P}^*(\cdot) = \mathbb{P}(\cdot | \mathbf{X}, \mathbf{Y})$ denote the conditional probability. In the following theorem, we establish the validity of the proposed bootstrap procedure.

Theorem 3.2. *Let Assumptions 3.1–3.3 hold. Assume that*

$$\Lambda_{\max}\|\Theta\|_\infty + \Delta_{\max} + \Delta_\sigma = o\left(\frac{1}{\log d}\right). \quad (3.7)$$

Then, under conditions of Theorem 3.1, we have

$$\sup_{x>0} \left| \mathbb{P} \left(\sqrt{n}\|\tilde{\beta}_\lambda - \beta^*\|_\infty \leq x \right) - \mathbb{P}^* \left(\hat{L} \leq x \right) \right| \xrightarrow{\mathbb{P}} 0.$$

Remark 8. Following the same de-biasing procedure as given in (3.2), we are also able to construct entrywise [[Javanmard and Montanari, 2014](#)] and groupwise [[Zhang and Cheng, 2017](#), [Dezeure et al., 2017](#)] simultaneous confidence intervals for β^* . For each $1 \leq j \leq d$, a $(1 - \alpha)$ -confidence interval for β_j^* is given by

$$\mathcal{CI}_\alpha(\beta_j^*) = \left\{ \tilde{\beta}_{j,\lambda} - \hat{\sigma} z_{1-\alpha/2} \sqrt{\frac{\hat{\Theta}_{jj}}{n}}, \tilde{\beta}_{j,\lambda} + \hat{\sigma} z_{1-\alpha/2} \sqrt{\frac{\hat{\Theta}_{jj}}{n}} \right\},$$

where $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -th quantile of standard normal distribution. For simultaneous groupwise inference of β^* , let G be a subset of $\{1, \dots, d\}$ of interest and consider testing the hypotheses

$$H_{0,G} : \beta_j^* = \beta_j^\circ \text{ for all } j \in G \text{ versus } H_{1,G} : \beta_j^* \neq \beta_j^\circ \text{ for some } j \in G.$$

In particular, when $\beta_j^\circ = 0$ for all $j \in G$, this reduces to testing the significance of a group of parameters. We obtain that the asymptotic distribution of $\max_{j \in G} \sqrt{n} |\tilde{\beta}_{j,\lambda} - \beta_j^*|$ converges to the distribution of $\max_{j \in G} |Z_j|$ by leveraging the Gaussian approximation. The remaining steps follow directly by conducting the Gaussian multiplier bootstrap. \square

4 Is Sparse Linear Model Adequate?

Sparse linear regression, which serves as the backbone of high dimensional statistics, has been widely applied in many areas of science, engineering, and social sciences. However, its adequacy has never been validated. This section focuses on testing the adequacy of the sparse linear model.

4.1 Main Results

As mentioned in introduction, the proposed model (1.5) contains the sparse linear regression model as a special case. Thus, we consider testing the hypotheses

$$H_0 : Y = \mathbf{x}^\top \boldsymbol{\beta}^* + \varepsilon \text{ versus } H_1 : Y = \mathbf{f}^\top \boldsymbol{\varphi}^* + \mathbf{x}^\top \boldsymbol{\beta}^* + \varepsilon, \quad (4.1)$$

which is equivalent to test whether $\boldsymbol{\varphi}^* = \boldsymbol{\gamma}^* - \mathbf{B}^\top \boldsymbol{\beta}^* = 0$. Since \mathbf{B} is an unknown dense matrix, simultaneously testing this linear equation will suffer from the curse of dimensionality.

On the other hand, for any set $\mathcal{S} \subset [d]$ with $\mathcal{S}_* \subset \mathcal{S}$, we have $\mathbf{B}_\mathcal{S}^\top \boldsymbol{\beta}_\mathcal{S}^* = \mathbf{B}^\top \boldsymbol{\beta}^*$. Hence, it suffices to compare the following two linear models in reduced dimension:

$$H_0 : Y = \mathbf{x}_\mathcal{S}^\top \boldsymbol{\beta}_\mathcal{S}^* + \varepsilon \text{ versus } H_1 : Y = \mathbf{f}^\top \boldsymbol{\varphi}^* + \mathbf{x}_\mathcal{S}^\top \boldsymbol{\beta}_\mathcal{S}^* + \varepsilon. \quad (4.2)$$

This hinges applying a sure screening method to reduce the dimensionality. There exist several methods which lead to the sure screening property. Among those, the commonly used one is the marginal

screening method [Fan and Lv, 2008, Fan and Song, 2010, Zhu et al., 2011, Li et al., 2012, Liu et al., 2014, Barut et al., 2016, Chu et al., 2016, Wang and Leng, 2016].

We propose an ANOVA-type test for (4.1) with two stages. In the first stage, the data set is split into two data sets $(\mathbf{Y}^{(1)}, \mathbf{X}^{(1)})$ and $(\mathbf{Y}^{(2)}, \mathbf{X}^{(2)})$, with sample sizes m and $n - m$, respectively. We use $(\mathbf{Y}^{(1)}, \mathbf{X}^{(1)})$ to screen variables. Let $\hat{\mathcal{S}}_1$ denote the set of variables selected. In the second stage, we leverage the selected $\hat{\mathcal{S}}_1$ and remaining data $(\mathbf{Y}^{(2)}, \mathbf{X}^{(2)})$ to perform hypothesis testing based on the ANOVA-type test statistic for low-dimensional model (4.2) with \mathcal{S} replaced by $\hat{\mathcal{S}}_1$. As the first step is based on marginal screening and is relatively crude, the sample size m is relatively small in comparing with the second step. We impose a general assumption on the set $\hat{\mathcal{S}}_1$.

Assumption 4.1 (Sure screening property). There exists a $s_n > 0$ such that

$$\mathbb{P} \left(|\hat{\mathcal{S}}_1| \leq s_n \text{ and } \mathcal{S}_* \subset \hat{\mathcal{S}}_1 \right) \rightarrow 1, \text{ as } n \rightarrow \infty.$$

A simple procedure that satisfies the above assumption is the follow factor-adjusted marginal screening based on the data $(\mathbf{Y}^{(1)}, \mathbf{X}^{(1)})$.

1. Estimation. Compute the latent factor estimator $\hat{\mathbf{F}}^{(1)}$, idiosyncratic component $\hat{\mathbf{U}}^{(1)}$ based on $\mathbf{X}^{(1)}$, and $\tilde{\mathbf{Y}}^{(1)} = (\mathbf{I}_m - \hat{\mathbf{F}}^{(1)}(\hat{\mathbf{F}}^{(1)\top}\hat{\mathbf{F}}^{(1)})^{-1}\hat{\mathbf{F}}^{(1)\top})\mathbf{Y}^{(1)}$.
2. Marginal regression. Compute the least square estimate $\hat{\beta}_{\ell, M} = \hat{\mathbf{U}}_{\ell}^{(1)\top}\tilde{\mathbf{Y}}^{(1)} / (\hat{\mathbf{U}}_{\ell}^{(1)\top}\hat{\mathbf{U}}_{\ell}^{(1)})$ for each $1 \leq \ell \leq d$.
3. Screening. Let $\hat{\mathcal{S}}_1 := \hat{\mathcal{S}}_{\phi} = \{\ell \in [d] : |\hat{\beta}_{\ell, M}| > \phi\}$ for some prescribed $\phi > 0$.

Here $\hat{\mathbf{U}}_{\ell}^{(1)} \in \mathbb{R}^d$ stands for the ℓ -th column of the matrix $\hat{\mathbf{U}}^{(1)}$. We next provide a sufficient condition for the Assumption 4.1 to hold.

Proposition 4.1. *Assume that $m = o(d \log d)$ and there exists some positive constant $\bar{c} < 1$ such that*

$$\phi < \frac{1}{1 + \bar{c}} \min_{\ell \in [d]} \frac{\Sigma_{\ell}^{\top} \beta^*}{\Sigma_{\ell\ell}} \text{ and } \frac{\mathcal{V}_{m,d}}{m} \|\varphi^*\|_2 + \|\beta^*\|_1 \frac{\log d}{m} + \|\beta^*\|_2 \sqrt{\frac{\log d}{m}} = o(\phi). \quad (4.3)$$

Here Σ_{ℓ} denotes the ℓ -th column of Σ . Then, under the Assumptions 2.1–2.5, we have

$$\mathbb{P} \left(\mathcal{S}_* \subset \hat{\mathcal{S}}_{\phi} \right) \rightarrow 1, \text{ as } m \rightarrow \infty.$$

Furthermore, we assume that $\min_{\ell \in \mathcal{S}_*} |\beta_{\ell,M}^*| := \Sigma_{\ell}^{\top} \beta^* / \Sigma_{\ell\ell} \geq c_* m^{-\kappa}$ for some positive constant $\kappa < 1/2$. Then for any $\phi = c_{\diamond} m^{-\kappa}$ with $c_{\diamond} \leq c_*/(1 + \bar{c})$, we have

$$\mathbb{P} \left\{ |\hat{\mathcal{S}}_{\phi}| \leq \frac{c_{\diamond}^2 m^{2\kappa} \|\Sigma \beta^*\|_2^2}{\lambda_{\min}^2(\Sigma)(1 - \bar{c})^2} \right\} \rightarrow 1 \text{ as } m \rightarrow \infty.$$

Remark 9. From the conclusion of Proposition 4.1, we obtain sure screening property by using our first data set with sample size $m = n^{\alpha}$ for some $\alpha < 1$ as long as the signal satisfies $\min_{\ell \in \mathcal{S}_*} |\beta_{\ell,M}^*| \geq c_* m^{-\kappa}$. Thus, the size of the remaining data set for constructing the test statistic in our second step is $n - n^{\alpha} \approx n$. It is worth to note that this does not lose any efficiency in terms of the asymptotic power in our hypothesis test when n goes to infinity. \square

Remark 10. Fan et al. [2020a] proposed a similar sure screening estimator which is a special case of our Proposition 4.1 with $\varphi^* = \gamma^* - \mathbf{B}^{\top} \beta^* = 0$. Moreover, we also provide an upper bound for the number of selected variables whereas Fan et al. [2020a] only provided a sufficient condition for the sure screening property. \square

Next, we proceed to the second stage of our hypothesis testing. In this step, we construct an ANOVA test statistic for (4.2) with \mathcal{S} replaced by $\hat{\mathcal{S}}_1$, which is given by

$$Q_n^{(2)} = \left\| \left(\mathbf{I}_{n-m} - \mathbf{P}_{\mathbf{X}_{\hat{\mathcal{S}}_1}^{(2)}} \right) \mathbf{Y}^{(2)} \right\|_2^2 - \left\| \left(\mathbf{I}_{n-m} - \mathbf{P}_{\hat{\mathbf{F}}^{(2)}} - \mathbf{P}_{\hat{\mathbf{U}}_{\hat{\mathcal{S}}_1}^{(2)}} \right) \mathbf{Y}^{(2)} \right\|_2^2. \quad (4.4)$$

We then summarize our results on the asymptotic behaviors of $Q_n^{(2)}$ in the following Theorem 4.2.

Theorem 4.2. Let Assumptions 2.1–2.5 and Assumption 4.1 hold with

$$s_n \left(\frac{\log d}{n} + \frac{1}{d} \right) \rightarrow 0 \text{ and } \Delta_{\sigma} \rightarrow 0.$$

We obtain

$$\sup_{x>0} \left| \mathbb{P} (Q_n^{(2)} \leq x \hat{\sigma}^2 | H_0) - \mathbb{P} (\chi_K^2 \leq x) \right| \rightarrow 0, \text{ as } n \rightarrow \infty.$$

Theorem 4.2 yields a level α test for (4.1) with critical region $\{Q_n^{(2)} > \hat{\sigma}^2 \chi_{K,1-\alpha}^2\}$, where $\chi_{K,1-\alpha}^2$ is the $(1 - \alpha)$ -th quantile of χ_K^2 -distribution.

Remark 11. Under stronger conditions such as irrepresentable condition [Zhao and Yu, 2006] or RIP condition [Candes and Tao, 2007], the $\hat{\mathcal{S}}$ achieved by certain explicit regularization [Zhao and Yu, 2006, Fan and Lv, 2011, Shi et al., 2019, Fan et al., 2020a] or implicit regularization accompanied with early stopping and signal truncation [Zhao et al., 2019, Fan et al., 2021c] enjoys variable selection consistency $\mathbb{P}(\hat{\mathcal{S}} = \mathcal{S}_*) \rightarrow 1$. In this scenario, we take the test statistic as

$$Q_n = \left\| \left(\mathbf{P}_{\hat{\mathbf{F}}} + \mathbf{P}_{\hat{\mathbf{U}}_{\hat{\mathcal{S}}}} - \mathbf{P}_{\mathbf{X}_{\hat{\mathcal{S}}}} \right) \mathbf{Y} \right\|_2^2$$

without using sample splitting. Under Assumptions 2.1–2.5, we obtain

$$\sup_{x>0} \left| \mathbb{P}(Q_n \leq x\hat{\sigma}^2 | H_0) - \mathbb{P}(\chi_K^2 \leq x) \right| \rightarrow 0, \quad (4.5)$$

by following similar proof idea with Theorem 4.2. \square

We now present the power of the test statistic (4.4).

Theorem 4.3. Define

$$\mathcal{D}(\alpha, \theta) = \left\{ \boldsymbol{\varphi} \in \mathbb{R}^K : \frac{n\|\boldsymbol{\varphi}\|^2}{1 + Ks_n\|\mathbf{B}\|_{\max}^2/\lambda_{\min}(\boldsymbol{\Sigma})} \geq \sigma^2(2 + \delta)(\chi_{K,1-\alpha}^2 + \chi_{K,1-\theta}^2) \right\},$$

where $\delta > 0$ is some constant, s_n is the size of selected set from the first stage and K is the number of factors. Assume that

$$\|\boldsymbol{\varphi}^*\|_2 \left(\sqrt{n/d} + 1/\sqrt{n} \right) \rightarrow 0. \quad (4.6)$$

Then, under the conditions of Theorem 4.2, we have

$$\inf_{\boldsymbol{\varphi}^* \in \mathcal{D}(\alpha, \theta)} \mathbb{P}(\psi_\alpha = 1 | H_1) \geq 1 - \theta.$$

Remark 12. Dataset with multiple types are now frequently collected for a common set of experimental subjects. This new data structure is also called multimodal data. It is worth to mention, the above hypothesis test can be further extended to test the adequacy of multi-modal sparse linear regression model [Li and Li, 2021]. To be more specific, we consider the hypothesis test as follows:

$$H_0 : \mathbf{Y} = \sum_{i=1}^L \mathbf{X}_i \boldsymbol{\beta}_i^* + \sum_{i=L+1}^M \mathbf{X}_i \boldsymbol{\beta}_i^* + \mathcal{E} \quad \text{versus} \quad H_1 : \mathbf{Y} = \sum_{i=1}^L (\mathbf{F}_i \boldsymbol{\gamma}_i^* + \mathbf{U}_i \boldsymbol{\beta}_i^*) + \sum_{i=L+1}^M \mathbf{X}_i \boldsymbol{\beta}_i^* + \mathcal{E}.$$

We aim at simultaneously testing whether the sparse regression is adequate for the any given L modals. Here $\mathbf{X}_i \in \mathbb{R}^{n \times d_i}$ is generated from the i -th modal, and possesses its own factor structure $\mathbf{X}_i = \mathbf{F}_i \mathbf{B}_i^\top + \mathbf{U}_i, i \in [M]$. Interested readers are referred to Appendix D.4 for more details. \square

5 Numerical Studies

5.1 Accuracy of Estimation

For data generation, we let number of factors $K = 2$, dimension of covariate $d = 1000$, $\gamma^* = (0.5, 0.5)$, the first $s = 3$ entries of β^* be 0.5 and remaining $d - s$ entries be 0. Throughout this subsection, we generate every entry of \mathbf{F}, \mathbf{U} from the standard Gaussian distribution and let every entry of \mathbf{B} be generated from the uniform distribution $\text{Unif}(-1, 1)$. We choose the noise distribution of ε given in model (2.1) from (i) standard Gaussian, (ii) uniform, and (iii) t_3 distribution respectively.

Distributions (i) and (ii) have sub-Gaussian tails. For these two cases, we select sample size n so that $s\sqrt{\log d/n}$ takes uniform grids in $[0.15, 0.5]$. Then we generate n response variables from model (2.1) and estimate our parameters via (2.3). The results are shown as the red lines in Figure 1. They lend further support to our theoretical findings given in section 2 as the statistical rates there are upper bounded by $O(s\sqrt{\log d/n})$. Moreover, we also show the estimation results by using Lasso directly on measurements (\mathbf{X}, \mathbf{Y}) . Results are shown as the blue lines given in the first two figures in Figure 1. Using Lasso directly on (\mathbf{X}, \mathbf{Y}) leads to much worse results due in part to the inadequacy of the model. In addition, as shown in [Fan et al. \[2020a\]](#), even when the sparse regression model is correct, we still have better estimation accuracy using factor adjusted regression.

Distribution (iii) has only the bounded second moment. Likewise, we select corresponding number of observations n so that $(s + K)\sqrt{\log d/n}$ takes uniform grids in $[0.4, 0.7]$. The reduced sample sizes help reduce the computation cost on the regularized adaptive Huber estimation using cross-validation to choose the parameter ω . We compare the results for the robust estimator (2.6) with that of the factor adjusted regression (2.3). The results are shown as the red and blue lines in part (c) of Figure 1 respectively. They provide stark evidence that it is necessary to conduct the robust version of factor

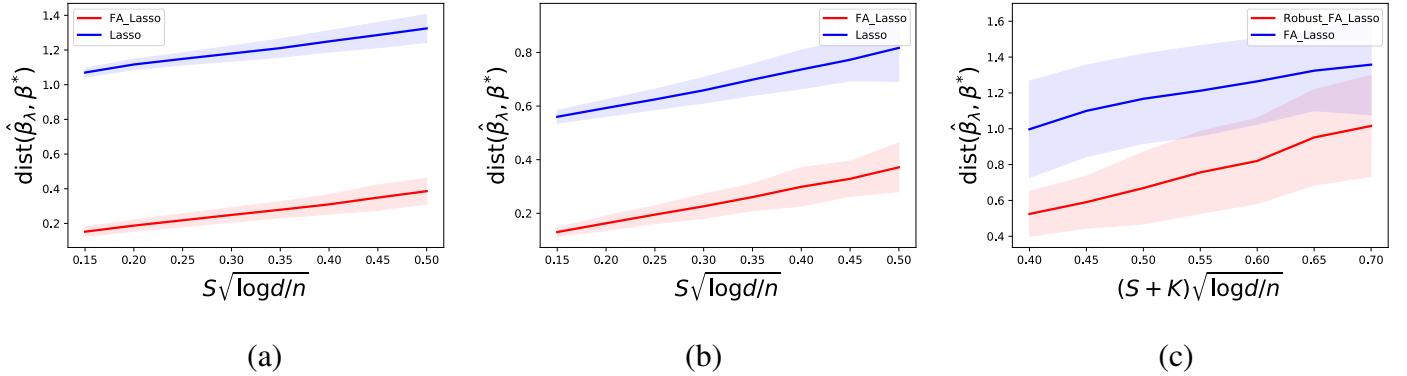


Figure 1: Accuracy for $\hat{\beta}_\lambda$ with $\text{dist}(\hat{\beta}_\lambda, \beta^*) := \|\hat{\beta}_\lambda - \beta^*\|_1$ based on 500 replications. The light color regions indicate the standard errors across the simulation. Figure (a) and (b) depict the estimation results of model (1.4) with noise ϵ following the standard Gaussian and uniform distributions respectively. In (a) and (b), the red lines denote the estimation results using the method (2.3) (labeled as FA_Lasso in the figure) and the blue lines represent the results using Lasso with data (\mathbf{X}, \mathbf{Y}) . In Figure (c), the noise ϵ follows t_3 -distribution. The red line in (c) represents the result of robust factor adjusted regression (Robust_FA_Lasso) via adaptive Huber estimation together with ℓ_1 -penalty given in (2.6) and the blue line represents the result achieved by using FA_Lasso.

adjusted regression (2.6) when noises have heavy tails.

5.2 Adequacy of Factor Regression

Data Generation Processes. We choose $n = 200$, $K = 2$ and d either 200 or 500 and the matrix $\mathbf{X} = \mathbf{F}\mathbf{B}^\top + \mathbf{U}$ using the following two models with entries of \mathbf{B} generated from $\text{Unif}(-1, 1)$.

1. We generate every row of $\mathbf{F} \in \mathbb{R}^{n \times K}$, $\mathbf{U} \in \mathbb{R}^{n \times d}$ from $\mathcal{N}(\mathbf{0}, \mathbf{I}_K)$ and $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ respectively.
2. We let the t -th row $\mathbf{f}_t \in \mathbb{R}^K$ of $\mathbf{F} \in \mathbb{R}^{n \times K}$ follow $\mathbf{f}_t = \Phi \mathbf{f}_{t-1} + \boldsymbol{\xi}_t$ where $\Phi \in \mathbb{R}^{K \times K}$ with $\Phi_{i,j} = 0.5^{|i-j|+1}$, $i, j \in [K]$. In addition, $\{\boldsymbol{\xi}_t\}_{t \geq 1}$ are drawn independently from $\mathcal{N}(\mathbf{0}, \mathbf{I}_K)$. We generate every row of \mathbf{U} from $\mathcal{N}(\mathbf{0}, \Sigma)$ where $\Sigma_{i,j} = 0.6^{|i-j|}$, $i, j \in [d]$.

The response vector follows $\mathbf{Y} = \mathbf{F}\boldsymbol{\gamma}^* + \mathbf{U}\boldsymbol{\beta}^* + \mathcal{E}$ in (4.1) with every entry of $\mathcal{E} \in \mathbb{R}^n$ being generated independently from either from $\mathcal{N}(0, 0.5^2)$ or uniform distribution $\text{Unif}(-\sqrt{3}/2, \sqrt{3}/2)$.

We set $\boldsymbol{\gamma}^* = (0.5, 0.5)$ and $\boldsymbol{\beta}^* = (w, w, w, 0, \dots, 0)$, where $w \geq 0$. When $w = 0$, the null hypothesis $\mathbf{Y} = \mathbf{F}\boldsymbol{\gamma}^* + \mathcal{E}$ holds and the simulation results correspond to the size of the test. Otherwise, they correspond to the power of the test.

Implementation. We summarize the details of the proposed test

1. Estimate factor $\hat{\mathbf{F}}$, loading matrix $\hat{\mathbf{B}}$, and noise $\hat{\mathbf{U}}$ as in section 2.
2. For the given $\hat{\mathbf{U}}$, we estimate $\hat{\boldsymbol{\Theta}}$ by using node-wise regression [van de Geer et al., 2014].
3. Use Lasso method to estimate $\hat{\boldsymbol{\beta}}_\lambda$ based on data $(\hat{\mathbf{U}}, \tilde{\mathbf{Y}})$ with λ chosen by cross-validation.
4. Construct a de-biased $\tilde{\boldsymbol{\beta}}_\lambda = \hat{\boldsymbol{\beta}}_\lambda + n^{-1}\hat{\boldsymbol{\Theta}}\hat{\mathbf{U}}^\top(\mathbf{Y} - \hat{\mathbf{U}}\hat{\boldsymbol{\beta}}_\lambda)$.
5. Obtain $\hat{L} = 1/\sqrt{n}\|\hat{\boldsymbol{\Theta}}\hat{\mathbf{U}}^\top\tilde{\mathcal{E}}\|_\infty$ with every entry of $\tilde{\mathcal{E}}$ simulated from $N(0, 1)$ and repeat this computation 1000 times to obtain the upper- α quantile $\hat{c}_{1-\alpha}$ of \hat{L} .
6. Compute the test result $\xi_\alpha := \mathbb{I}\{\sqrt{n}\|\tilde{\boldsymbol{\beta}}_\lambda\|_\infty/\hat{\sigma} > \hat{c}_{1-\alpha}\}$ where $\hat{\sigma}$ is estimated by using refitted cross-validation [Fan et al., 2012].

The size and the power are then calculated based on 2000 simulations with $\alpha = 0.05$.

Results. For $n = 200$, $K = 2$, $d \in \{200, 500\}$ and all $w \in \{0, 0.05, 0.10, 0.15, 0.20\}$, we generate the data from each model and compute the testing results based on 2000 simulations with $\alpha = 0.05$. The results are depicted in the Table 1. The column named Gaussian(i), $i \in \{1, 2\}$ represents the simulation results under model i with Gaussian noise. Similar labels applied to the uniform noise distribution.

Table 1 reveals that our test gives approximately the right size (subject to simulation error; see the rows with $w = 0$). This is consistent with our theoretical findings given in section 3. In addition, when $0 < w < 0.2$, the power of our test increases rapidly to 1 which reveals the efficiency of our test statistic.

5.3 Adequacy of Sparse Regression

This subsection provides finite-sample validations for the results in section 4. We take the number of data used for screening $m = \lceil n^{0.8} \rceil$, use Iterative Sure Independence Screening method [Fan and

		Gaussian (1)	Gaussian (2)	Uniform (1)	Uniform (2)
$p = 200$	$w = 0$	0.044	0.047	0.046	0.048
	$w = 0.05$	0.067	0.119	0.065	0.108
	$w = 0.10$	0.326	0.714	0.311	0.653
	$w = 0.15$	0.859	0.989	0.854	0.984
	$w = 0.20$	0.998	1.000	0.996	1.000
$p = 500$	$w = 0$	0.043	0.040	0.048	0.436
	$w = 0.05$	0.067	0.080	0.059	0.071
	$w = 0.10$	0.253	0.632	0.237	0.563
	$w = 0.15$	0.787	0.974	0.780	0.962
	$w = 0.20$	0.993	1.000	0.987	1.000

Table 1: Simulation results of section 3 under different regimes.

[Lv, 2008, Saldana and Feng, 2018, Zhang et al., 2019] to select $\hat{\mathcal{S}}_1$ and apply the refitted cross-validation [Fan et al., 2012] to estimate σ^2 . The size and the power of the test are computed based on 2000 simulations.

Data Generation Processes. We let $n = 250$, $K = 3$ and d be either 250 or 600. The noises ε are i.i.d from $N(0, 0.5^2)$ or $\text{Unif}(-\sqrt{3}/2, \sqrt{3}/2)$. The covariate $\mathbf{X} \in \mathbb{R}^{n \times d}$ follows the factor model $\mathbf{X} = \mathbf{F}\mathbf{B}^\top + \mathbf{U}$. We generate \mathbf{F} , \mathbf{U} and \mathbf{B} in the same way as those in section 5.2. In addition, the response variable follows $\mathbf{Y} = \mathbf{F}\boldsymbol{\varphi}^* + \mathbf{X}\boldsymbol{\beta}^* + \mathcal{E}$ in (4.1) with $\boldsymbol{\beta}^* = (0.8, 0.8, 0.8, 0.8, 0, \dots, 0)$ and $\boldsymbol{\varphi}^* = v \cdot \mathbf{1}_{K \times 1}$ for several different values of $v \geq 0$. The case $v = 0$ corresponds to the null hypothesis and it is designed to test the validity of the size.

Results. For $n = 250$, $K = 3$, $d \in \{250, 600\}$ and $v \in \{0, 0.04, 0.08, 0.12, 0.16\}$, we implement the proposed method for every model in section 5.2. The simulation results are depicted in Table 2. The column named Gaussian (or uniform) (i) , $i \in \{1, 2\}$ represents the results under model i with Gaussian (or uniform) noise mentioned in section 5.2. When $v = 0$, the null hypothesis holds, our Type-I error is approximately 0.05 which matches with the theoretical value. In addition, when we increase the size of v from $v = 0.04$ to $v = 0.16$, the power of our test statistic increases sharply to 1, which reveals its efficiency.

		Gaussian (1)	Gaussian (2)	Uniform (1)	Uniform (2)
$p = 250$	$v = 0$	0.051	0.054	0.056	0.053
	$v = 0.04$	0.215	0.278	0.233	0.286
	$v = 0.08$	0.659	0.740	0.655	0.750
	$v = 0.12$	0.965	0.993	0.965	0.996
	$v = 0.16$	1.000	1.000	1.000	1.000
$p = 600$	$v = 0$	0.051	0.052	0.050	0.052
	$v = 0.04$	0.208	0.362	0.197	0.353
	$v = 0.08$	0.624	0.802	0.604	0.785
	$v = 0.12$	0.941	0.994	0.934	0.999
	$v = 0.16$	1.000	1.000	0.999	1.000

Table 2: Simulation results of section 4 under different regimes.

We next discuss the necessity of using sample splitting. Suppose we do not split samples and use the whole dataset to do sure screening and construct the test statistic. This will result in the high correlation between the selected set $\hat{\mathcal{S}}$ and covariates when $\hat{\mathcal{S}}$ is not a consistent estimator of \mathcal{S}_* . In this case, the asymptotic behavior of our test statistic is hard to capture. To demonstrate this point, we simulate the null distribution of the test statistic constructed without using sample splitting and compare it with the asymptotic distribution (χ_K^2) via the quantile-quantile plot in Figure 2. Figure 2 reveals that the test statistic constructed without using sample splitting has heavier right tail than that of the χ_K^2 distribution. The sizes of the test are much larger than the results in Table 2 when $v = 0$.

5.4 Empirical Applications

In this section, we use a macroeconomic dataset named FRED-MD [McCracken and Ng, 2016] to illustrate the performance of our factor augmented regression model (FARM) and investigate whether the latent factor regression model and sparse linear model are adequate.

There are 134 monthly U.S. macroeconomic variables in this dataset. As they measure certain aspects of economic health, these variables are driven by latent factors and hence correlated. They can be well explained by a few principal components. In our study, we pick out two variables named

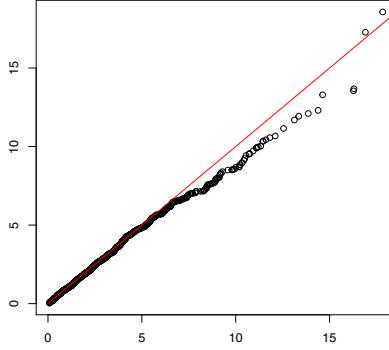


Figure 2: Quantiles of the χ_K^2 distribution against those of the test statistic without sample splitting. The x-axis represents the quantiles of the test statistic whereas y-axis is the quantiles of χ_K^2 distribution.

'HOUSTNE' and 'GS5' as our responses respectively and let the remaining variables be the covariates. Here 'HOUSTNE' represents the housing starts in the northeast region and 'GS5' denotes the 5-year treasury rate.

There exist significant structural breaks for many variables around the year of financial crisis in 2008 which makes our data non-stationary even after performing the suggested transformations. Thus, we analyze the dataset in two separate time periods independently. Specifically, we study the monthly data collected from February 1992 to October 2007 and from August 2010 to February 2020 respectively after examining the missingness and stationarity of the data.

To illustrate the performance of our proposed FARM against the sparse linear model and latent factor regression model, we first analyze the prediction results achieved by using these models. For every given time period and model, we perform the prediction by using the moving window approach with window size 90 months. Indexing the panel data from 1 for each of the two time periods, for all $t > 90$, we use the 90 previous measurements $\{(\mathbf{x}_{t-90}, Y_{t-90}), \dots, (\mathbf{x}_{t-1}, Y_{t-1})\}$ to train a model (FARM, sparse linear regression model, or latent factor regression model) and output a prediction \hat{Y}_t as well as the in-sample average $\bar{Y}_t := \frac{1}{90} \sum_{i=t-90}^{t-1} Y_i$. We measure the prediction accuracy by using

out-of-sample R^2 :

$$R^2 = 1 - \frac{\sum_{t=91}^T (Y_t - \hat{Y}_t)^2}{\sum_{t=91}^T (Y_t - \bar{Y}_t)^2},$$

where T denotes the number of total data points in a given time period. Table 3 presents the out-of-sample R^2 obtained by the aforementioned three models in the two time periods for predicting 'HOUSTNE' and 'GS5'. Their detailed predictions are depicted in Figure 3 and Figure 4 respectively. Both out-of-sample R^2 and predictions depicted in Figures 3 and 4 show that FARM outperforms both latent factor regression and sparse linear regression models.

Time period	Data	FARM	SP.Linear	LA.Factor
02.1992-10.2007	HOUSTNE	0.632	0.584	0.428
	GS5	0.705	0.631	0.122
08.2010-02.2020	HOUSTNE	0.694	0.450	0.219
	GS5	0.657	0.540	0.233

Table 3: Out-of-sample R^2 for predicting 'HOUSTNE' and 'GS5' data using different models in different time periods. LA.Factor and SP.Linear stand for latent factor regression and sparse linear regression respectively.

We next conduct the hypothesis testing on the adequacy of latent factor regression and sparse linear regression respectively by using FARM as the alternative model. As computing the bootstrap estimate of the null distribution is expensive for testing the adequacy of the factor model, we only conduct the hypothesis testing using the data in the entire two subperiods: 02.1992-10.2007 and 08.2010-02.2020. The P-values for the tests are given in Table 4. Taking the significant level 0.05, the hypothesis testing results indicate that the latent factor regression is not adequate for all four different settings in Table 4. As for sparse linear regression, it is accepted only for studying 'HOUSTNE' in the time period 02.1992-10.2007. These results match well with our prediction results.

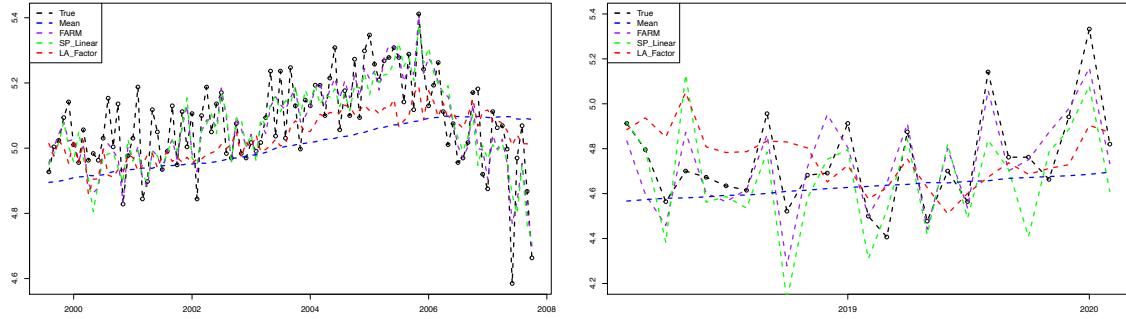


Figure 3: Out-of-sample prediction results for 'HOUSTNE' data in time periods: August 1999 to October 2007 (left panel) and February 2018 to February 2020 (right panel). The black dash line represents the true observed values, and the blue, purple, green and red dash lines represent the predictions made by using in-sample mean (moving average with a window of 90), FARM, sparse linear model and latent factor regression model, respectively.

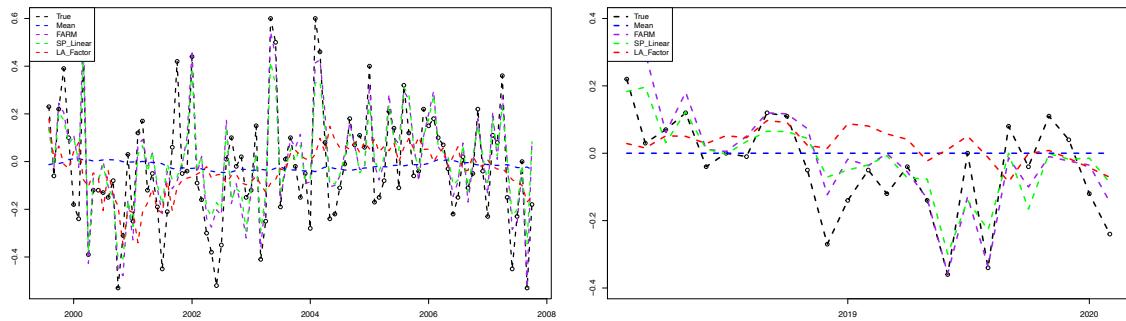


Figure 4: Out-of-sample prediction results for 'GS5' data in time periods: August 1999 to October 2007 (left panel) and February 2018 to February 2020 (right panel). The same captions as those in Figure 3 are used.

Time period	Data	LA_factor	SP_Linear
02.1992-10.2007	HOUSTNE	$2.00 \cdot 10^{-3}$	$5.75 \cdot 10^{-1}$
	GS5	$2.50 \cdot 10^{-3}$	$8.73 \cdot 10^{-3}$
08.2010-02.2020	HOUSTNE	$< 10^{-3}$	$2.00 \cdot 10^{-2}$
	GS5	$1.70 \cdot 10^{-2}$	$2.94 \cdot 10^{-2}$

Table 4: p -values for testing the adequacy of the latent factor regression and sparse linear regression models to explain 'HOUSTNE' and 'GS5' data in two different time periods. The LA_Factor and SP_Linear have the same meaning as those in Table 3.

6 Conclusion

In this paper, we propose a model named Factor Augmented sparse linear Regression Model (FARM), which contains the latent factor regression and the sparse linear regression as our special cases. The model expands the space spanned by covariates into useful directions and hence use additional information beyond the linear space spanned by the predictors. We provide theoretical guarantees for our model estimation under the existence of light-tailed and heavy-tailed noises respectively. In addition, we leverage our model as the alternative one to test the sufficiency of the latent factor regression model and sparse regression model. We believe that the study is among the first of this kind. The practical performance of our model estimation and our constructed test statistics are proven by extensive simulation studies including both synthetic data and real data. Moreover, it is worth to mention that our model and methodology can be extended to more general supervised learning problems such as non-parametric regression, quantile regression, regression and classification trees, support vector machines, among others where the factor augmentation idea is always useful.

References

S. C. Ahn and A. R. Horenstein. Eigenvalue ratio test for the number of factors. *Econometrica*, 81(3):1203–1227, 2013.

M. Avella-Medina, H. S. Battey, J. Fan, and Q. Li. Robust estimation of high-dimensional covariance and precision matrices. *Biometrika*, 105(2):271–284, 2018.

J. Bai. Inferential theory for factor models of large dimensions. *Econometrica*, 71(1):135–171, 2003.

J. Bai and K. Li. Statistical analysis of factor models of high dimension. *Ann. Statist.*, 40(1):436–465, 2012.

J. Bai and S. Ng. Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221, 2002.

J. Bai and S. Ng. Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica*, 74(4):1133–1150, 2006.

J. Bai and S. Ng. Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2):304–317, 2008. ISSN 0304-4076.

E. Bair, T. Hastie, D. Paul, and R. Tibshirani. Prediction by supervised principal components. *Journal of the American Statistical Association*, 101(473):119–137, 2006.

E. Barut, J. Fan, and A. Verhasselt. Conditional sure independence screening. *J. Amer. Statist. Assoc.*, 111(515):1266–1277, 2016.

A. Belloni and V. Chernozhukov. ℓ_1 -penalized quantile regression in high-dimensional sparse models. *Ann. Statist.*, 39(1):82–130, 2011.

X. Bing, F. Bunea, and M. Wegkamp. Inference in latent factor regression with clusterable features. *arXiv:1905.12696*, 2019.

X. Bing, F. Bunea, S. Strimas-Mackey, and M. Wegkamp. Prediction under latent factor regression: Adaptive pcr, interpolating predictors and beyond. *Journal of Machine Learning Research*, 22(177):1–50, 2021.

F. Bunea, S. Strimas-Mackey, and M. Wegkamp. Interpolating predictors in high-dimensional factor regression. *arXiv:2002.02525*, 2020.

T. Cai, W. Liu, and X. Luo. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *J. Amer. Statist. Assoc.*, 106(494):594–607, 2011.

E. Candes and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Statist.*, 35(6):2313–2351, 2007.

V. Chernozhukov, D. Chetverikov, and K. Kato. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *Ann. Statist.*, 41(6):2786–2819, 2013.

V. Chernozhukov, D. Chetverikov, and K. Kato. Central limit theorems and bootstrap in high dimensions. *Ann. Probab.*, 45(4):2309–2352, 2017.

V. Chernozhukov, D. Chetverikov, and Y. Koike. Nearly optimal central limit theorem and bootstrap approximations in high dimensions. *arXiv preprint arXiv:2012.09513*, 2020.

W. Chu, R. Li, and M. Reimherr. Feature screening for time-varying coefficient models with ultrahigh dimensional longitudinal data. *Ann. Appl. Stat.*, 10(2):596, 2016.

R. Dezeure, P. Bühlmann, and C.-H. Zhang. High-dimensional simultaneous inference with the bootstrap. *Test*, 26(4):685–719, 2017.

B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004.

J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96(456):1348–1360, 2001.

J. Fan and Y. Liao. Learning latent factors from diversified projections and its applications to over-estimated and weak factors. *Journal of the American Statistical Association*, 0(0):1–16, 2020.

J. Fan and J. Lv. Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 70(5):849–911, 2008.

J. Fan and J. Lv. Nonconcave penalized likelihood with NP-dimensionality. *IEEE Trans. Inform. Theory*, 57(8):5467–5484, 2011.

J. Fan and R. Song. Sure independence screening in generalized linear models with NP-dimensionality. *Ann. Statist.*, 38(6):3567–3604, 2010.

J. Fan, S. Guo, and N. Hao. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 74(1):37–65, 2012.

J. Fan, Y. Liao, and M. Mincheva. Large covariance estimation by thresholding principal orthogonal complements. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 75(4):603–680, 2013. With 33 discussions by 57 authors and a reply by Fan, Liao and Mincheva.

J. Fan, Y. Fan, and E. Barut. Adaptive robust variable selection. *Ann. Statist.*, 42(1):324, 2014.

J. Fan, Q. Li, and Y. Wang. Estimation of high dimensional mean regression in the absence of symmetry and light tail assumptions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 79(1):247, 2017a.

J. Fan, L. Xue, and j. Yao. Sufficient forecasting using factor models. *Journal of Econometrics*, 201(2):292–306, 2017b. ISSN 0304-4076.

J. Fan, Y. Ke, and K. Wang. Factor-adjusted regularized model selection. *J. Econometrics*, 216(1):71–85, 2020a.

J. Fan, R. Li, C.-H. Zhang, and H. Zou. Statistical foundations of data science. 2020b.

J. Fan, J. Guo, and S. Zheng. Estimating number of factors by adjusted eigenvalues thresholding. *Journal of the American Statistical Association*, pages 1–10, 2021a.

J. Fan, R. Masini, and M. C. Medeiros. Bridging factor and sparse models. *arXiv:2102.11341*, 2021b.

J. Fan, Z. Yang, and M. Yu. Understanding implicit regularization in over-parameterized single index model. *arXiv:2007.08322v3*, 2021c.

M. Hernan and J. Robins. *Causal Inference*. Chapman & Hall/CRC Monographs on Statistics & Applied Probab. CRC Press LLC, 2019. ISBN 9781420076165.

H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.

G. Imbens and D. Rubin. Causal inference for statistics, social, and biomedical sciences: An introduction. 2015.

A. Javanmard and A. Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.*, 15:2869–2909, 2014.

I. T. Jolliffe. A note on the use of principal components in regression. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 31(3):300–303, 1982.

C. Lam and Q. Yao. Factor modeling for high-dimensional time series: Inference for the number of factors. *The Annals of Statistics*, 40(2):694 – 726, 2012.

G. Li, H. Peng, J. Zhang, and L. Zhu. Robust rank correlation based screening. *Ann. Statist.*, 40(3):1846–1877, 2012.

Q. Li and L. Li. Integrative factor regression and its inference for multimodal data analysis. *J. Amer. Statist. Assoc.*, 113(521):1–15, 2021.

Q. Li, G. Cheng, J. Fan, and Y. Wang. Embracing the blessing of dimensionality in factor models. *J. Amer. Statist. Assoc.*, 113(521):380–389, 2018.

J. Liu, R. Li, and R. Wu. Feature selection for varying coefficient models with ultrahigh-dimensional covariates. *J. Amer. Statist. Assoc.*, 109(505):266–274, 2014.

P.-L. Loh and M. J. Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *Ann. Statist.*, 40(3):1637–1664, 2012.

M. W. McCracken and S. Ng. Fred-md: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589, 2016.

R. Nickl and S. Van De Geer. Confidence sets in sparse regression. *Ann. Statist.*, 41(6):2852–2876, 2013.

B. Peng, L. Wang, and Y. Wu. An error bound for L_1 -norm support vector machine coefficients in ultra-high dimension. *J. Mach. Learn. Res.*, 17(1):8279–8304, 2016.

D. F. Saldana and Y. Feng. Sis: An r package for sure independence screening in ultrahigh-dimensional statistical models. *Journal of Statistical Software*, 83(2):1–25, 2018.

C. Shi, R. Song, Z. Chen, and R. Li. Linear hypothesis testing for high dimensional generalized linear models. *Ann. Statist.*, 47(5):2671–2703, 2019.

J. H. Stock and M. W. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460):1167–1179, 2002.

Q. Sun, W.-X. Zhou, and J. Fan. Adaptive Huber regression. *Journal of the American Statistical Association*, 115(529):254–265, 2020.

T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.

R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.

S. Van de Geer. High-dimensional generalized linear models and the lasso. *Ann. Statist.*, 36(2):614–645, 2008.

S. van de Geer, P. Bühlmann, Y. Ritov, and R. Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.*, 42(3):1166–1202, 2014.

W. Wang and J. Fan. Asymptotics of empirical eigenstructure for high dimensional spiked covariance. *Ann. Statist.*, 45(3):1342–1374, 2017.

X. Wang and C. Leng. High dimensional ordinary least squares projection for screening variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, pages 589–611, 2016.

G. Yu and J. Bien. Estimating the error variance in a high-dimensional linear model. *Biometrika*, 106(3):533–546, 2019.

C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942, 2010.

C.-H. Zhang and S. S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 76(1):217–242, 2014.

N. Zhang, W. Jiang, and Y. Lan. On the sure screening properties of iteratively sure independence screening algorithms. *arXiv:1812.01367*, 2019.

X. Zhang and G. Cheng. Simultaneous inference for high-dimensional linear models. *J. Amer. Statist. Assoc.*, 112(518):757–768, 2017.

X. Zhang, Y. Wu, L. Wang, and R. Li. Variable selection for support vector machines in moderately high dimensions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 78(1):53, 2016.

P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.

P. Zhao, Y. Yang, and Q.-C. He. Implicit regularization via hadamard product over-parametrization in high-dimensional linear regression. *arXiv:1903.09367*, 2019.

L.-P. Zhu, L. Li, R. Li, and L.-X. Zhu. Model-free feature screening for ultrahigh-dimensional data. *J. Amer. Statist. Assoc.*, 106(496):1464–1475, 2011.

H. Zou. The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429, 2006.