Published in final edited form as:

Stat Sin. 2018 May ; 28: 2633–2655. doi:10.5705/ss.202017.0027.

Partial Consistency with Sparse Incidental Parameters

Jianqing Fan¹, Runlong Tang², Xiaofeng Shi¹

¹Princeton University

²Johns Hopkins University

Abstract

Penalized estimation principle is fundamental to high-dimensional problems. In the literature, it has been extensively and successfully applied to various models with only structural parameters. As a contrast, in this paper, we first apply this penalization principle to a linear regression model with a finite-dimensional vector of structural parameters and a high-dimensional vector of sparse incidental parameters. For the estimators of the structural parameters, we derive their consistency and asymptotic normality, which reveals an oracle property. However, the penalized estimators for the incidental parameters possess only partial selection consistency but not consistency. This is an interesting partial consistency phenomenon: the structural parameters are consistently estimated while the incidental ones cannot. For the structural parameters, also considered is an alternative two-step penalized estimator, which has fewer possible asymptotic distributions and thus is more suitable for statistical inferences. We further extend the methods and results to the case where the dimension of the structural parameter vector diverges with but slower than the sample size. A data-driven approach for selecting a penalty regularization parameter is provided. The finite-sample performance of the penalized estimators for the structural parameters is evaluated by simulations and a real data set is analyzed.

Keywords

Structural Parameters; Sparse Incidental Parameters; Penalized Estimation; Partial Consistency; Oracle Property; Two-Step Estimation; Confidence Intervals

1. Introduction

Since the pioneering papers by Tibshirani (1996) and Fan and Li (2001), the penalized estimation methodology for exploiting sparsity has been studied extensively. For example, Zhao and Yu (2006) provides an almost necessary and sufficient condition, namely Irrepresentable Condition, for the LASSO estimator to be strong sign consistent. Fan and Lv (2011) shows that an oracle property holds for the folded concave penalized estimator with ultrahigh dimensionality. For an overview on this topic, see Fan and Lv (2010).

All the aforementioned papers consider models only with the so-called structural parameters, which are related to every data point. By contrast, we consider in this paper another type of models where there are not only the structural parameters but also the so-called incidental parameters, each of which is related to only one data point. Specifically, suppose data $\{X_i, Y_i\}_{i=1}^n$ are from the following linear model:

$$Y_i = \mu_i^* + X_i^T \beta^* + \epsilon_i, \quad (1.1)$$

where the vector of *incidental* parameters $\mu^* = (\mu_1^*, \dots, \mu_n^*)^T$ is sparse, the vector of *structural* parameters $\beta^* = (\beta_1^*, \dots, \beta_d^*)^T$ is of main interest, and for each i , X_i is a d -dimensional covariate vector, and ϵ_i is a random error. Let $\nu = (\mu^{*T}, \beta^{*T})^T$. Then, in model (1.1), a different data point (X_i, Y_i) depends on a different subset of ν , that is, μ_i^* and β^* .

Model (1.1) arises as a working model for estimation from Fan et al. (2012b), which considers a large-scale hypothesis testing problem under arbitrary dependence of test statistics. By Principal Factor Approximation, a method proposed by Fan et al. (2012b), the dependent test statistics $\mathbf{Z} = (Z_1, \dots, Z_p)^T \sim N(\mu, \Sigma)$ can be decomposed as $Z_i = \mu_i + b_i^T \mathbf{W} + K_i$, where $\mu = (\mu_1, \dots, \mu_p)^T$ and b_i is the i th row of the first k unstandardized principal components, denoted by $\mathbf{B}$, of Σ and $\mathbf{K} = (K_1, \dots, K_p)^T \sim N(0, \mathbf{A})$ with $\mathbf{A} = \mathbf{I} - \mathbf{B}\mathbf{B}^T$. The common factor $\mathbf{W}$ drives the dependence among the test statistics. This realized but unobserved factor is critical for False Discovery Proportion (FDP) estimation and power improvements by removing the common factor $\{b_i^T \mathbf{W}\}$ from the test statistics. Hence, an important goal is to estimate $\mathbf{W}$ with give $\{b_i\}_{i=1}^n$. In many applications on large-scale hypothesis testing, the parameters $\{\mu_i\}_{i=1}^p$ are sparse. For example, genome-wide association studies show that the expression level of gene CCT8 is highly related to the phenotype of Down Syndrome. It is of interest to test the association between each of millions of SNP's and the CCT8 gene expression level. In the framework of Fan et al. (2012b), each μ_i stands for such an association. That is, if $\mu_i = 0$, the i th SNP is not associated with the CCT8 gene expression level; otherwise, it is associated. Since most of the SNP's are not associated the CCT8 gene expression level, it is reasonable to assume $\{\mu_i\}_{i=1}^p$ are sparse. Replacing $Z_i, \mu_i, b_i, \mathbf{W}, k, p$ and K_i with $Y_i, \mu_i^*, X_i, \beta^*, d, n$ and ϵ_i respectively, we obtain model (1.1) formally. It is of interest to study this model independently with simplifications.

Although model (1.1) emerges from a critical component of estimating FDP in Fan et al. (2012b), it stands with its own interest. For example, in some applications, there are only few signals (nonzero μ_i^* 's) and what is interesting is to learn about β^* , which reflects the relationship between the covariates and response. For another example, those few nonzero

μ_i^* 's might be some measurement or recording errors of the responses $\{Y_i\}$. In these cases, model (1.1) is suitable for modeling data with contaminated responses and a method producing a reliable estimator for β^* is essentially a robust replacement for ordinary least squares estimate, which is sensitive to outliers.

Several models with structural and incidental parameters have first been studied in a seminal paper by Neyman and Scott (1948), which points out the inconsistency of the maximum likelihood estimators (MLE) of structural parameters in the presence of a large number of incidental parameters and provides a modified MLE. However, their method does not work for model (1.1) due to no exploration of the sparsity of incidental parameters. Kiefer and Wolfowitz (1956) shows the consistency of the MLE of the structural parameters when the incidental parameters are assumed to be from a common distribution. That is, they eliminate the essential high-dimensional issue of the incidental parameters by randomizing them. In contrast, this paper considers deterministic incidental parameters and handles the high-dimensional issue by penalization with a sparsity assumption. Basu (1977) considers the elimination of nuisance parameters via marginalizing and conditioning methods and Moreira (2009) solves the incidental parameter problem with an invariance principle. For a review of the incidental parameter problems in statistics and economics, see Lancaster (2000).

Without loss of generality, suppose the first s incidental parameters $\{\mu_i^*\}_{i=1}^s$ are nonvanishing and the remaining are zero. Then, model (1.1) can be written in a matrix form as $Y = X\nu + \epsilon$, where

$$X = \begin{pmatrix} I_s & X_{1,s}^T & 0 \\ 0 & X_{s+1,n}^T & I_{n-s} \end{pmatrix},$$

$X_{i,j}^T = (X_i, X_{i+1}, \dots, X_j)^T$, I_k is a $k \times k$ identity matrix, $\mathbf{0}$ is a generic block of zeros and $\nu = (\mu_1^*, \dots, \mu_s^*, \beta^T, \mu_{s+1}^*, \dots, \mu_n^*)^T$. Although this is a sparse high-dimensional problem, the matrix X does not satisfy the sufficient conditions of the theoretical results in Zhao and Yu (2006) and Fan and Lv (2011) due to the inconsistency of the estimation of the incidental parameters in ν and the penalty should not be simply placed on all parameters. For details, see Supplement C.

In this paper, we investigate a penalized estimator of (μ^*, β^*) defined by

$$(\hat{\mu}, \hat{\beta}) = \underset{(\mu, \beta) \in \mathbb{R}^{n+d}}{\operatorname{argmin}} \sum_{i=1}^n (Y_i - \mu_i - X_i^T \beta)^2 + \sum_{i=1}^n p_\lambda(|\mu_i|), \quad (1.2)$$

where p_λ is a penalty function with a regularization parameter λ . Since only the incidental parameters are sparse, the penalty is imposed on them. An iterative algorithm is proposed to numerically compute the estimators. The estimator $\hat{\beta}$ possesses consistency, asymptotic

normality and an oracle property. On the other hand, the nonvanishing elements of μ^* cannot be consistently estimated even if β^* were known. So, there is a partial consistency phenomenon.

Penalized estimation (1.2) is a one-step method. For the estimation of β^* , We also propose a two-step method whose first step is designed to eliminate the influence of the data with large incidental parameters. The estimator $\tilde{\beta}$ from the two-step method has fewer possible asymptotic distributions than $\hat{\beta}$ and thus is more suitable for constructing confidence regions for β^* . It is asymptotically equivalent to the one-step estimator $\hat{\beta}$ when the sizes of the nonzero incidental parameters are small enough, that is, when the incidental parameters are really sparse. Also, the two-step method improves the convergence rate and efficiency over the one-step method for challenging situations where large nonzero incidental parameters increase the asymptotic covariance or even reduce the convergence rate for the one-step method.

The rest of the paper is organized as follows. In Section 2, the model and penalized estimation method are formally introduced and the corresponding penalized estimators are characterized. In Section 3, asymptotic properties of the penalized estimators are derived; a penalized two-step estimator is proposed and its theoretical properties are obtained; we also provide a data-driven approach for selecting the regularization parameter. In Section 4, we consider the case where the number of covariates grows with but slower than the sample size. In Section 5, we present simulation results and analyze a read data set. Section 6 concludes this paper with a discussion and all the proofs and some theoretical results are relegated to the appendix and supplements.

2 Model and Method

The matrix form of model (1.1) is given by

$$Y = \mu^* + X\beta^* + \epsilon, \quad (2.1)$$

where $Y = (Y_1, Y_2, \dots, Y_n)^T$, $X = (X_1, X_2, \dots, X_n)^T$, and $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_n)^T$. The covariates $\{X_i\}_{i=1}^n$ are independent and identically distributed (i.i.d.) copies of $X_0 \in \mathbb{R}^d$, which is a random vector with mean zero and a covariance matrix $\Sigma_X > 0$. They are independent of the random errors $\{\epsilon_i\}$, which are i.i.d. copies of ϵ_0 , which is a random variable with mean zero and variance $\sigma^2 > 0$. Denote $a_n \ll b_n$ and $a_n \gg b_n$ if $a_n = o(b_n)$ and $b_n = o(a_n)$, respectively. There is an assumption on the covariates and random errors.

Assumption (A): There exist positive sequences $\kappa_n \ll \sqrt{n}$, $\gamma_n \ll \sqrt{n}$ such that

$$P\left(\max_{1 \leq i \leq n} \|X_i\|_2 > \kappa_n\right) \rightarrow 0 \text{ and } P\left(\max_{1 \leq i \leq n} |\epsilon_i| > \gamma_n\right) \rightarrow 0, \text{ as } n \rightarrow \infty, \quad (2.2)$$

where $\|\cdot\|_2$ stands for the L_2 norm of $\mathbb{R}^d$.

Suppose there are three types of incidental parameters in model (1.1) or (2.1): for simplicity on the indexes, the first s_1 incidental parameters $\{\mu_i^*\}_{i=1}^{s_1}$ are large in the sense that $|\mu_i^*| \gg \max\{\kappa_n, \gamma_n\}$ for $1 \leq i \leq s_1$; the next s_2 ones $\{\mu_i^*\}_{i=s_1+1}^s$ are nonzero and bounded by γ_n with $s = s_1 + s_2$; the last $n - s$ ones $\{\mu_i^*\}_{i=s+1}^n$ are zero. Note that it is unknown to us which μ_i^* 's are large, bounded or zero. The sparsity of μ^* is understood by $s_1 + s_2 \ll n$, i.e. $s_1 + s_2 = o(n)$. Denote the vectors of the three types of incidental parameters $\mu_1^*, \mu_2^*,$ and μ_3^* , respectively.

The penalized estimation (1.2) can be written as

$$(\hat{\mu}, \hat{\beta}) = \underset{(\mu, \beta)}{\operatorname{argmin}} L(\mu, \beta), \quad L(\mu, \beta) = \left\| Y - \mu - X\beta \right\|_2^2 + \sum_{i=1}^n p_\lambda(|\mu_i|). \quad (2.3)$$

The penalty function p_λ can be the soft (i.e. L_1 or LASSO), hard, SCAD or a general folded concave penalty function (Fan and Li, 2001). For simplicity, we next consider only the soft penalty function, that is, $p_\lambda(|\mu_i|) = 2\lambda|\mu_i|$. The cases with the hard and SCAD penalties can be considered in a similar way.

By Lemma D.1 in Supplement D, A necessary and sufficient condition for $(\hat{\mu}, \hat{\beta})$ to be a minimizer of $L(\mu, \beta)$ is that $\hat{\beta} = (X^T X)^{-1} X^T (Y - \hat{\mu})$, $Y_i - \hat{\mu}_i - X_i^T \hat{\beta} = \lambda \operatorname{Sign}(\hat{\mu}_i)$ for $i \in \hat{I}_0^c$ and $|Y_i - X_i^T \hat{\beta}| \leq \lambda$ for $i \in \hat{I}_0$, where $\operatorname{Sign}(\cdot)$ is a sign function and $\hat{I}_0 = \{1 \leq i \leq n: \hat{\mu}_i = 0\}$.

Numerically, the special structure of $L(\mu, \beta)$ suggests a marginal decent algorithm for the minimization problem in (2.3), which iteratively computes $\mu^{(k)} = \underset{\mu \in \mathbb{R}^n}{\operatorname{argmin}} L(\mu, \beta^{(k-1)})$ and

$\beta^{(k)} = \underset{\beta \in \mathbb{R}^d}{\operatorname{argmin}} L(\mu^{(k)}, \beta)$ until convergence. The advantage of this algorithm is that there exist

analytic solutions to the above two minimization problems. They are respectively the soft-threshold estimators with residuals $\{Y_i - X_i^T \beta^{(k-1)}\}$ and ordinary least-squares estimator with responses $Y - \mu^{(k)}$. In this section and the next, the number of the covariates d is assumed to be a *fixed* integer. A case where d diverging to infinity will be considered in Section 4. For the case where d is finite, we make the following assumption on λ .

Assumption (B): The regularization parameter λ satisfies

$$\kappa_n \ll \lambda, \alpha \gamma_n \leq \lambda, \text{ and } \lambda \ll \min\{\mu^*, \sqrt{n}\}, \quad (2.4)$$

where k_n and γ_n are defined in (2.2), α is a constant greater than 2, and $\mu^\star = \min_{1 \leq i \leq s_1} |\mu_i^\star|$.

For simplicity, abbreviate “with probability going to one” to “wpg1”. A stopping rule for the above algorithm is based on the successive difference $\|\beta^{(k+1)} - \beta^{(k)}\|_2$. By Proposition D.2 in Supplement wpg1, the iterative algorithm stops at the the second iteration, given the initial estimator is bounded wpg1.

Suppose $\{\beta^{(k)}\}$ has a theoretical limit $\beta^{(\infty)}$, corresponding to which, there is a limit estimator $\mu^{(\infty)}$. Then, $(\mu^{(\infty)}, \beta^{(\infty)})$ is a solution of the following system of nonlinear equations

$$\beta = (X^T X)^{-1} X^T (Y - \mu), \quad (2.5)$$

and, with soft-threshold applied to each component, it follows

$$\mu = (|Y - X^T \beta| - \lambda)_+ \text{Sign}(Y - X^T \beta), \quad (2.6)$$

where $(\cdot)_+$ returns the maximum value of the input and zero. By Lemma D.3 in Supplement D, a necessary and sufficient condition for $(\hat{\mu}, \hat{\beta})$ to be a minimizer of $L(\mu, \beta)$ is that it is a solution to equations (2.5) and (2.6). Hence, $(\mu^{(\infty)}, \beta^{(\infty)})$ is a minimizer of $L(\mu, \beta)$ and can also be denoted as $(\hat{\mu}, \hat{\beta})$.

Note that $\hat{\beta}$ is also the minimizer of the profiled loss function $\tilde{L}(\beta) = L(\mu(\beta), \beta)$, where $\mu(\beta)$ as a function of β is given by (2.6). Interestingly, this profiled loss function is a criterion function equipped with the famous Huber loss function (see Huber (1964) and Huber (1973)). Specifically, the profiled loss function can be expressed as $\tilde{L}(\beta) = \sum_{i=1}^n \rho(Y_i - X_i^T \beta)$ where $\rho(x) = x^2 I(|x| \leq \lambda) + (2\lambda x - \lambda^2) I(|x| > \lambda)$ is *exactly* the Huber loss function, which is optimal in a minimax sense. This equivalence between the penalized estimation and Huber’s robust estimation indicates that the penalization principle is versatile and can naturally produce an important loss function in robust statistics. This equivalence also provides a formal endorsement of the least absolute deviation robust regression (LAD) in Fan et al. (2012b) and indicates that it is better to use all data with LAD regression rather than 90% of them. It is worthwhile to note that the penalized estimation is only formally equal to the Huber’s. Our model (2.1) considers *deterministic* sparse incidental parameters $\mu_i^\star$ ’s, while the model in Huber’s works assumes *random* contamination as in Kiefer and Wolfowitz (1956). Recently, there appear a few papers on robust regression in high-dimensional settings, see, for example, Chen et al. (2010), Lambert-Lacroix and Zwald (2011), ? and Bean et al. (2012). Portnoy and He (2000) provide a high level review of literature on robust statistics.

From the equations (2.5) and (2.6), $\hat{\beta}$ is a solution to

$$\varphi_n(\beta) = 0, \text{ where } \varphi_n(\beta) = \beta - (X^T X)^{-1} X^T (Y - \mu(\beta)). \quad (2.7)$$

In general, this is a Z-estimation problem. The following theoretical analysis is based on this characterization of $\hat{\beta}$.

At the end of this section, we provide for further analysis some notations and an expansion of $\varphi_n(\beta)$. Let $S = \sum_{i=1}^n X_i X_i^T$, $\mathbb{S}_S = \sum_{i \in S} X_i X_i^T$, $\mathbb{S}_S^\mu = \sum_{i \in S} X_i \mu_i^*$, $S_S^\epsilon = \sum_{i \in S} X_i \epsilon_i$, $\mathcal{S} = \sum_{i=1}^n X_i$ and $\mathcal{S}_S = \sum_{i \in S} X_i$, where S is a subset of $\{1, 2, \dots, n\}$. It is straightforward to show

$$\begin{aligned} \varphi_n(\beta) = & \left(\mathbb{S}_{S_{10}} + \mathbb{S}_{S_{11}} + \mathbb{S}_{S_{12}} \right) (\beta - \beta^*) - \left(\mathbb{S}_{S_{11}}^\mu + \mathbb{S}_{S_{12}}^\mu \right) \\ & - \left(\mathbb{S}_{S_{10}}^\epsilon + \mathbb{S}_{S_{11}}^\epsilon + \mathbb{S}_{S_{12}}^\epsilon \right) - \lambda \left(\mathcal{S}_{S_{20}} + \mathcal{S}_{S_{21}} + \mathcal{S}_{S_{22}} - \mathcal{S}_{S_{30}} - \mathcal{S}_{S_{31}} - \mathcal{S}_{S_{32}} \right), \end{aligned} \quad (2.8)$$

where the index sets $S_{10} = \{s+1 \leq i \leq n: |X_i^T(\beta^* - \beta) + \epsilon_i| \leq \lambda\}$, $S_{11} = \{1 \leq i \leq s_1: |\mu_i^* + X_i^T(\beta^* - \beta) + \epsilon_i| \leq \lambda\}$ and $S_{12} = \{s_1+1 \leq i \leq s: |\mu_i^* + X_i^T(\beta^* - \beta) + \epsilon_i| \leq \lambda\}$; S_{20} , S_{21} and S_{22} are defined similarly except that the absolute operation is omitted and “ $\leq$ ” is replaced by “ $>$ ”; S_{30} , S_{31} and S_{32} are defined similarly with S_{20} , S_{21} and S_{22} except that “ $>$ ” is replaced by “ $<$ ”. Note that all these index sets depend on β .

3 Asymptotic Properties

In this section, we consider the asymptotic properties of the penalized estimators $\hat{\beta}$ and $\hat{\mu}$. Assumption (A), together with Assumption (B), enables the penalized estimation method to distinguish the large incidental parameters from others, and thus simplifies the asymptotic properties of the index sets S_{ij} 's in (2.8) in the sense that they become independent of β wpg1. Denote a hypercube of β^* by $B_C(\beta^*) = \{\beta \in \mathbb{R}^d: |\beta_j - \beta_j^*| \leq C, 1 \leq j \leq d\}$ with a constant $C > 0$.

Lemma 3.1 (On Index Sets S_{ij} 's). *Under Assumptions (A) and (B), for every $C > 0$ and every $\beta \in B_C(\beta^*)$, wpg1, $S_{10} = S_{10}^*$, $S_{11} = \emptyset$, $S_{12} = S_{12}^*$, $S_{20} = \emptyset$, $S_{21} = S_{21}^*$, $S_{22} = \emptyset$, $S_{30} = \emptyset$, $S_{31} = S_{31}^*$ and $S_{32} = \emptyset$, where the limit index sets $S_{10}^* = \{s+1, s+2, \dots, n\}$, $S_{12}^* = \{s_1+1, s_1+2, \dots, s\}$, $S_{21}^* = \{1 \leq i \leq s_1: \mu_i^* > 0\}$ and $S_{31}^* = \{1 \leq i \leq s_1: \mu_i^* < 0\}$.*

By Lemma 3.1, wpgl, the solution $\hat{\beta}$ to (2.7) has an analytic expression:

$$\hat{\beta} = \beta^* + \left(\mathbb{S}_{S_{10}^*} + \mathbb{S}_{S_{12}^*} \right)^{-1} \left[\mathbb{S}_{S_{12}^*}^{\mu} + \left(\mathbb{S}_{S_{10}^*}^{\epsilon} + \mathbb{S}_{S_{12}^*}^{\epsilon} \right) + \lambda \left(\mathcal{S}_{S_{21}^*} - \mathcal{S}_{S_{31}^*} \right) \right], \quad (3.1)$$

from which, we derive asymptotic properties of $\hat{\beta}$. Some analysis needs the following assumption.

Assumption (C): There exists some constant $\delta > 0$ such that $\mathbb{E}\|X_0\|_2^{2+\delta} < \infty$ and

$$\|\mu_2^*\|_2 / \|\mu_2^*\|_{2+\delta} \text{ diverges to infinity, where } \|\mu_2^*\|_{2+\delta} = \left(\sum_{i=s_1+1}^s |\mu_i^*|^{2+\delta} \right)^{1/(2+\delta)}.$$

The following result shows the existence of a unique consistent estimator of β^* .

Theorem 3.2 (Existence and Consistency of $\hat{\beta}$). *Under Assumptions (A) and (B), if either $s_2 = o(n/(k_n \gamma_n))$ or Assumption (C) holds, then, for every fixed $C > 0$, wpgl, there exists a unique estimator $\hat{\beta}_n \in B_C(\beta^*)$ such that $\psi_n(\hat{\beta}_n) = 0$ and $\hat{\beta}_n \xrightarrow{P} \beta^*$.*

In Theorem 3.2, there are two different kinds of sufficient conditions: one is on s_2 , which is the size of bounded incidental parameters μ_2^* , and the other is Assumption (C), which is about the norms of μ_2^* . They come from different analysis approaches on the term $\mathbb{S}_{S_{12}^*}^{\mu}$ in

(3.1). One does not imply the other. For details, see Supplement E. Specially, if $s_2 = O(n^{\alpha_2})$

for some $\alpha_2 \in (0, 1)$ and $\kappa_n \gamma_n \ll n^{(1-\alpha_2)}$, then $\hat{\beta}$ is consistent by Theorem 3.2.

Next, we consider the asymptotic distributions of the consistent estimator $\hat{\beta}_n$ obtained in Theorem 3.2. Without loss of generality, we assume the sizes of index sets $S_{21}^* = \{1 \leq i \leq s_1 : \mu_i^* < 0\}$ and $S_{31}^* = \{1 \leq i \leq s_1 : \mu_i^* < 0\}$ are asymptotically equivalent to as_1 and $(1-\alpha)s_1$ with a constant $a \in (0, 1)$. Similar to Theorem 3.2, there are two different sets of conditions on μ_2^* corresponding to two different analysis approaches. Denote $\sim$ as the asymptotic equivalence and $D_n = \|\mu_2^*\|_2$.

Theorem 3.3 (Asymptotic Distributions on $\hat{\beta}_n$). *Under Assumptions (A) and (B), suppose $s_2 \ll \sqrt{n}/(\kappa_n \gamma_n)$ holds or Assumption (C) and $D_n^2/n = o(1)$ hold.*

1. If $s_1 \ll n/\lambda^2$, then $\sqrt{n}(\hat{\beta}_n - \beta^*) \xrightarrow{d} N(0, \sigma^2 \Sigma_X^{-1})$; [main case]
2. If $s_1 \sim bn/\lambda^2$, then $\sqrt{n}(\hat{\beta}_n - \beta^*) \xrightarrow{d} N(0, (b + \sigma^2) \Sigma_X^{-1})$, for every constant $b \in \mathbb{R}^+$;

3. If $s_1 \gg n/\lambda^2$, then $r_n(\hat{\beta}_n - \beta^*) \xrightarrow{d} N(0, \Sigma_X^{-1})$, where $r_n \sim n/(\lambda\sqrt{s_1})$.

When the incidental parameters are really sparse, the size s_1 of large incidental parameters is small and the size s_2 or the magnitude D_n of bounded incidental parameters is also small so that the conditions of case (1) tends to hold. This case is of most interest and we denote it as the *main case*. The other cases are presented to provide a relatively complete picture of the asymptotic distributions of $\hat{\beta}$. In fact, Theorem E.1 in Supplement E shows more possible asymptotic distributions. Note that the constant a does not appear in the limit distributions of Theorem 3.3 due to cancelation and that the sub $-\sqrt{n}$ consistency emerges in case (3) when s_1 is large, because for this case the impact of the large incidental parameters is too big to be handled efficiently by the penalized estimation. For case (2), in one direction, as $b \rightarrow 0$, its condition and limit distribution become those of case (1); in the other direction, as b increases, it approaches case (3). This boundary phenomenon was in spirit similar to that in Tang et al. (2012). Specially, if $\lambda \ll n^{\alpha_1}$, $\kappa_n \gamma_n \ll n^{\alpha_2}$, $s_1 \ll n^{1-\alpha_1}$ and $s_2 \ll n^{1/2-\alpha_2}$, for some $\alpha_1 \in (0,1)$ and $\alpha_2 \in (0,1/2)$, then $\sqrt{n}(\hat{\beta}_n - \beta^*) \xrightarrow{d} N(0, \sigma^2 \Sigma_X^{-1})$ by the main case of Theorem 3.3.

Remark 1 (An Oracle Property). Suppose an oracle tells the true μ^* . Then, with the adjusted responses $Y - \mu^*$, the oracle estimator of μ^* is given by $\hat{\beta}^{(O)} = (XX^T)^{-1}X^T(Y - \mu^*)$. The limiting distribution of $\sqrt{n}(\hat{\beta}_n^{(O)} - \beta^*)$ is $N(0, \sigma^2 \Sigma_X^{-1})$. Comparing this with the main case of Theorems 3.3, it follows that the penalized estimator $\hat{\beta}_n$ enjoys an oracle property.

Although mainly interested in the estimation of β^* , we also obtain the soft-threshold estimator $\hat{\mu}$ of μ^* : for each i ,

$$\hat{\mu}_i = \mu_i(\hat{\beta}) = \left(|Y_i - X_i^T \hat{\beta}| - \lambda \right)_+ \text{sgn}(Y_i - X_i^T \hat{\beta}). \quad (3.2)$$

Denote $\mathcal{E} = \{\hat{\mu}_i \neq 0, \text{ for } i = 1, 2, \dots, s_1; \text{ and } \hat{\mu}_i = 0, \text{ for } i = s_1 + 1, s_1 + 2, \dots, n\}$.

Theorem 3.4 (Partial Selection Consistency on $\hat{\mu}$). Under Assumptions (A) and (B), if $\hat{\beta} \xrightarrow{P} \beta^*$ $P(\mathcal{E}) \rightarrow 1$.

Theorem 3.4 shows that, wpg1, the indexes of μ_1^* and μ_3^* are estimated correctly, but those of μ_2^* wrongly. We call this a partial selection consistency phenomenon.

3.1 Two-Step Estimation

Theorems 3.3 shows that the penalized estimator $\hat{\beta}_n$ has multiple different limit distributions, which complicates the application of these theorems in practice. In addition, the convergence rate of $\hat{\beta}_n$ is less than the optimal rate $\sqrt{n}$ in the challenging cases where the impact of large incidental parameters is substantial. To address these issues, we propose the following two-

step estimation method: firstly, we apply the penalized estimation (2.3) and let $\hat{I}_0 = \{1 \leq i \leq n: \hat{\mu}_i = 0\}$; secondly, we define the two-step estimator as

$$\tilde{\beta} = \left(\mathbf{X}_{\hat{I}_0}^T \mathbf{X}_{\hat{I}_0} \right)^{-1} \mathbf{X}_{\hat{I}_0}^T \mathbf{Y}_{\hat{I}_0}, \quad (3.3)$$

Where $\mathbf{X}_{\hat{I}_0}$ consists of X_i 's whose indexes are in $\hat{I}_0$ and $\mathbf{Y}_{\hat{I}_0}$ consists of the corresponding Y_i 's. The following theorem shows that $\tilde{\beta}$ is consistent and Asymptotic Gaussian.

Theorem 3.5 (Consistency and Asymptotic Normality on $\tilde{\beta}$). *Suppose Assumptions (A) and (B) hold. If either $s_2 = o(n/(\kappa_n \gamma_n))$ or Assumption (C) holds, then $\tilde{\beta} \xrightarrow{P} \beta^*$. If either $s_2 = O(\sqrt{n}/(\kappa_n \gamma_n))$ holds or Assumptions (C) and $D_n^2/n = o(1)$ hold, then $\sqrt{n}(\tilde{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 \Sigma_X^{-1})$.*

Comparing Theorem 3.5 with Theorem 3.3, we see that $\hat{\beta}$ has three possible asymptotic distributions but $\tilde{\beta}$ has only one since for $\tilde{\beta}$ the conditions on s_1 disappear. It is because the two-step method identifies and removes large incidental parameters by exploiting the partial selection consistency property of $\hat{\mu}$ in Theorem 3.4. Further, the two-step estimator improves the convergence rate to the optimal one over the one-step estimator for the challenging case with $s_1 \gg n/\lambda^2$. Because of these advantages, we suggest to use the two-step method to make statistical inferences.

When the incidental parameters are sparse in the sense that the size or the magnitude of the bounded incidental parameters are small, i.e. $s_2 = o(\sqrt{n}/(\kappa_n \gamma_n))$ or $D_n^2/n = o(1)$, it follows, by Theorem 3.5, $\sqrt{n}(\tilde{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 \Sigma_X^{-1})$, from which a confidence region with asymptotic confidence level $1 - \alpha$ is given by

$$\left\{ \beta \in \mathbb{R}^d: \sigma^{-1} \sqrt{n} \left\| \Sigma_X^{1/2} (\tilde{\beta} - \beta) \right\|_2 \leq q_\alpha(\chi_d) \right\}, \quad (3.4)$$

Where $q_\alpha(\chi_d)$ is the upper α -quantile of χ_d , the square root of the chi-squared distribution with degrees of freedom d . For each component β_j^* of β^* , an asymptotic $1 - \alpha$ confidence interval is given by

$$\left[\tilde{\beta}_j \pm n^{-1/2} \sigma \Sigma_X^{-1/2}(j, j) z_{\alpha/2} \right], \quad (3.5)$$

where $\Sigma_X^{-1/2}(j, j)$ is the square root of the (j, j) entry of Σ_X^{-1} and $z_{\alpha/2}$ is the upper $\alpha/2$ -quantile of $N(0,1)$. The confidence region (3.4) and interval (3.5) involve unknown parameters Σ_X and σ . They can be estimated by $\hat{\Sigma}_X = (1/n)X^T X$ and

$$\hat{\sigma} = \# \left(\hat{I}_0 \right)^{-1/2} \left\| Y_{\hat{I}_0} - X_{\hat{I}_0}^T \tilde{\beta} \right\|_2. \quad (3.6)$$

By Law of Large Numbers, $\hat{\Sigma}_X$ is consistent. By lemma E.3 in Supplement, $\hat{\sigma}$ is also consistent. Hence, after replacing Σ_X and σ in the confidence region (3.4) and interval (3.5) with $\hat{\Sigma}_X$ and $\hat{\sigma}$, the resulting confidence region and interval have the asymptotic confidence level $1 - \alpha$.

3.2 Theoretical and Data-Driven Regularization Parameters

By Assumption (B), the theoretical regularization parameter λ depends on k_n and γ_n , which are also crucial to the boundary conditions of the asymptotic properties of the penalized estimators $\hat{\beta}$ and $\tilde{\beta}$. Under Assumption (A), k_n and γ_n is determined by the distributions of X_0 and ϵ_0 , respectively. It is of interest to explicitly derive k_n and γ_n for some typical cases on the covariates and errors. When the covariates are bounded with $C_X > 0$ and the random errors follow $N(0, \sigma^2)$, let $\kappa_n = \sqrt{d}C_X$ and $\gamma_n = \sqrt{2\sigma^2 \log(n)}$. They satisfy (2.2) in Assumption (A) and the specification of λ (2.4) in Assumption (B) becomes $\alpha\sqrt{2\sigma^2 \log(n)} \leq \lambda \ll \min\{\mu^*, \sqrt{n}\}$. When X_0 and ϵ_0 follow $N(0, \Sigma_X)$ and $N(0, \sigma^2)$, respectively. Denote by σ_X^2 the maximum of diagonal elements of Σ_X . We can take $\kappa_n = \sqrt{2d\sigma_X^2 \log(n)}$ and $\gamma_n = \sqrt{2\sigma^2 \log(n)}$. Then, (2.4) becomes $\sqrt{\log(n)} \ll \lambda \ll \min\{\mu^*, \sqrt{n}\}$. A case on exponentially tailed random variables has all been considered in Supplement E.2.

Although the theoretical specification of λ guarantees desired asymptotic properties, a data-driven specification is of interest in practice. A popular way to specify λ is to use multi-fold cross-validation. The validation set, however, needs to be made as little contaminated as possible. We propose the following procedure to identify a data-driven regularization parameter:

Step 1: On the training and testing data sets.

1. Apply ordinary least squares (OLS) to all the data and obtain residuals $\hat{\epsilon}_i^{(OLS)} = Y_i - X_i^T \hat{\beta}^{(OLS)}$ for each i .
2. Identify the set of “pure” data corresponding to the n_{pure} smallest values in $\{|\hat{\epsilon}_i^{(OLS)}|\}$.
3. Compute the updated OLS estimator $\hat{\beta}^{(OLS,2)}$ with the “pure” data and obtain updated residuals $\{\hat{\epsilon}_i^{(OLS,2)}\}$ for each i .

4. Identify the updated “pure” data set corresponding to the n_{pure} smallest $\left\{ \left| \hat{\epsilon}_i^{(OLS, 2)} \right| \right\}$ and label the remaining as the “contaminated” data set.
5. Randomly select a subset from the updated “pure” data set as a testing set and merge the remaining “pure” data set and the “contaminated” one into a training set.

Step 2: On the range $[\lambda_L, \lambda_U]$ of the regularization parameter.

1. Compute the standard deviation $\hat{\sigma}_{pure}$ of the residuals of the “pure” data set.
2. Set $\lambda_L = \alpha_l \hat{\sigma}_{pure}$ and $\lambda_U = \alpha_u \hat{\sigma}_{pure}$ where $\alpha_l < \alpha_u$ are positive constants.

Step 3: On the data-driven regularization parameter.

1. For each grid point of λ in the interval $[\lambda_L, \lambda_U]$, apply a penalized method to the training set and obtain the estimator $\hat{\beta}_{\lambda, train}$ and the corresponding test
$$\hat{\sigma}_{\lambda, test}^2 = \sum_{\text{testing set}} (Y_i - X_i^T \hat{\beta}_{\lambda, train})^2.$$
2. Identify the data-driven regularization parameter λ_{opt} which minimizes $\hat{\sigma}_{\lambda, test}^2$ among the grid points.

This simple data-driven procedure can certainly be further improved. For example, In Step 1, the sub-steps 3 and 4 can be repeated more times to obtain a better “pure” data set. In Step two, the range for λ can also be obtained from quantiles of $\left\{ \left| \hat{\epsilon}_i^{(OLS, 2)} \right| \right\}$. We can also hybrid quantities based on $\hat{\sigma}_{pure}$ and quantiles of $\left\{ \left| \hat{\epsilon}_i^{(OLS, 2)} \right| \right\}$ to determine $[\lambda_L, \lambda_U]$.

The good performance of this data-driven regularization parameter will be demonstrated in Subsection 5.2.

4. Diverging number of structural parameters

In Sections 2 and 3, we have considered model (2.1) under the assumption that the number of covariates d is a fixed integer. However, when there are a moderate or large number of covariates, it is appropriate to assume that d diverges to infinity with the sample size. In this section, we consider model (2.1) with the assumption that $d \rightarrow \infty$ and $d \ll n$.

Since the number of covariates grows orderly slower than the sample size, we chose to continue use the penalized estimation (2.3) for (μ^*, β^*) and the penalized two-step estimation (3.3) for β^* . The corresponding estimators are still denoted as $(\hat{\mu}, \hat{\beta})$ and $\tilde{\beta}$, but we should keep it in mind that their dimensions diverge to infinity with n . The characterizations of $\hat{\beta}$ in Lemmas D.1 and D.3 are still valid since they are finite-sample results. The iteration algorithm also stops at the second iteration, which is shown by Proposition F.3 in Supplement F.

As before, it is critical to properly specify the regularization parameter λ for the case with a diverging number of covariates.

Assumption (B'): The regularization parameter λ satisfies

$$\sqrt{d}\kappa_n \ll \lambda, \alpha\gamma_n \leq \lambda, \text{ and } \lambda \ll \mu^*, \quad (4.1)$$

where k_n and γ_n are defined in (2.2) and $\alpha > 2$.

Comparing Assumption (B') with Assumption (B), the main difference in formation is that k_n is changed to $\sqrt{d}\kappa_n$. In fact, k_n in (4.1) also depends on d , which will be shown in Supplement E.2. This difference is caused by the assumption that d diverges to ∞ .

Lemma 4.1 (On Index Sets S_{ij} 's). *Under Assumptions (A) and (B'), the conclusion of Lemma 3.1 holds.*

Thus, wpg1, still valid is the crucial analytic expression of $\hat{\beta}$ (3.1), from which we derive its theoretical properties. They are similar to those in the previous section, with additional technical complexity caused by the diverging dimension d .

Denote $\|\cdot\|_{F,d} = d^{-1/2} \|\cdot\|_F$, where $\|\cdot\|_F$ is the Frobenius norm, and

$\kappa_X = d^{-1} \sum_{j=1}^d \left(\mathbb{E} [X_{0j}^4] \right)^{1/2}$ the average of the square root of the fourth marginal moments of X_0 . We make the following assumptions on Σ_X and κ_X .

Assumption (D): $\|\Sigma_X^{-1}\|_{F,d}$ is bounded.

Assumption (E): κ_X is bounded.

Theorem 4.2 (Existence and Consistency on $\hat{\beta}$). *Suppose Assumptions (A), (B'), (D) and (E) hold. If there exists r_d , a sequence of positive numbers depending on d , such that $d^3/n \rightarrow 0$, $(r_d d)^2/n \rightarrow 0$, $(r_d d)^2/n \rightarrow 0$, $s_1 = o(n/(r_d \sqrt{d}\kappa_n \lambda))$ and $s_2 = o(n/(r_d \sqrt{d}\kappa_n \gamma_n))$, then, for every fixed $C > 0$, wpg1, there exists a unique estimator $\hat{\beta} \in B_C(\beta^*)$ such that $\psi_n(\hat{\beta}) = 0$ and $r_d \|\hat{\beta} - \beta^*\|_2 \xrightarrow{P} 0$.*

Next, we consider the asymptotic distribution on $\hat{\beta}$. Since the dimension of $\hat{\beta}$ diverges to infinity, following Fan and Lv (2011), it is more appropriate to study its linear maps. Let A_n be a $q \times d$ matrix, where q is a fixed integer, $G_n = A_n A_n^T$ with the largest eigenvalue $\lambda_{\max}(G_n)$, and $G_{X,n} = A_n \Sigma_X^{-1} A_n^T$. Denote by $\lambda_{\min}(\Sigma_X)$ the smallest eigenvalue of Σ_X , $\sigma_{X,\max}^2 = \max_{1 \leq j \leq d} \text{Var}[X_{0j}]$, $\sigma_{X,\min}^2 = \min_{1 \leq j \leq d} \text{Var}[X_{0j}]$ and $\gamma_{X,\max} = \max_{1 \leq j \leq d} \mathbb{E}[X_{0j}]^3$. Abbreviate “with respect to” by “wrt”. We assume further

Assumption (D'): $\lambda_{\min}(\Sigma_X)$ is bounded away from zero, which implies Assumption (D).

Assumption (D''): $\|\Sigma_X\|_{F,d}$ is bounded.

Assumption (F): $\|A_n\|_F$ and $\lambda_{\max}(G_n)$ are bounded and $G_{X,n}$ converges to a $q \times q$ symmetric matrix G_X wrt $\|\cdot\|_F$.

Assumption (G): $\sigma_{X,\max} > 0$; $\sigma_{X,\max}$ and $\gamma_{X,\min}$ are bounded from above and $\sigma_{X,\min}$ is bounded away from zero.

Similar to the main case of Theorem 3.3, a properly scaled $\hat{\beta}_n$ is asymptotically Gaussian.

Theorem 4.3 (Asymptotic Distribution on $\hat{\beta}$). *Suppose Assumptions (A), (B'), (D'), (D''), (E), (F) and (G) hold. If $d^5 \log d = o(n)$, $s_1 = o(\sqrt{n}/(\lambda\sqrt{d}\kappa_n))$ and $s_2 = o(\sqrt{n}/(\sqrt{d}\kappa_n\gamma_n))$, then*

$$\sqrt{n}A_n(\hat{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 G_X).$$

The penalized estimator $\hat{\mu}$ obtained by (3.2) is partially consistent.

Theorem 4.4 (Partial Selection Consistency on $\hat{\mu}$). *Suppose Assumptions (A) and (B') hold and $\hat{\beta}$ is a consistent estimator of β^* wrt $r_d\|\cdot\|_2$. If $r_d \geq 1/\sqrt{d}$, then $P(\mathcal{E}) \rightarrow 1$.*

We can construct the penalized two-step estimator $\tilde{\beta}$ through (3.3) with $\hat{\mu}$. This two-step estimator is consistent by Theorem F.4 in Supplement F and its asymptotic distribution, as an extension of the main case in Theorem 3.5, is given by the following theorem.

Theorem 4.5 (Asymptotic Distribution on $\tilde{\beta}$). *Suppose all the assumptions and conditions of Theorem 4.3 hold except that the condition on s_1 is not required. Then*

$$\sqrt{n}A_n(\tilde{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 G_X).$$

From Theorems 4.3 and 4.5, Wald-type asymptotic confidence regions of β^* are available. For example, a confidence region based on $\tilde{\beta}$ with asymptotic confidence level $1 - \alpha$ is given by

$$\left\{ \beta \in \mathbb{R}^d : \sigma^{-1} \sqrt{n} \left\| G_{X,n}^{-1/2} A_n(\tilde{\beta} - \beta) \right\|_2 \leq q_\alpha(\chi_q) \right\}. \quad (4.2)$$

Since $G_{X,n}$ involves the unknown Σ_X , we estimate it by $\hat{G}_{X,n} = A_n \hat{\Sigma}_X^{-1} A_n^T$. On the other hand, σ is estimated by $\hat{\sigma}$ in (3.6) as before. After plugging $\hat{G}_{X,n}$ and $\hat{\sigma}$ into (4.2), we obtain

$$\left\{ \beta \in \mathbb{R}^d : \hat{\sigma}^{-1} \sqrt{n} \left\| \hat{G}_{X,n}^{-1/2} A_n(\tilde{\beta} - \beta) \right\|_2 \leq q_\alpha(\chi_q) \right\}. \quad (4.3)$$

By Lemma A.5 in the appendix, the consistency of $\hat{\sigma}$ is assured. Then, Theorem A.6 in the appendix guarantees the asymptotic validity of the confidence region (4.3).

5 Numerical Evaluations and Real Data Analysis

In this section, we evaluate the finite-sample performance of the penalized estimation through simulations and use it to analyze a real data set. The model for simulations is given by, for $i = 1, \dots, n$, $Y_i = \mu_i^* + X_{i,1}\beta_1^* + \dots + X_{i,50}\beta_{50}^* + \epsilon_i$. For simplicity, the deterministic sparse incidental parameters $\{\mu_i^*\}$ are generated as i.i.d. copies of μ : μ is 0, $U[-c, c]$ and $W(c + \text{Exp}(\tau))$ with probabilities p_0 , p_1 , and p_2 respectively, where $U[-c, c]$ is a uniform random variable in $[-c, c]$, W takes values 1 and -1 with probabilities p_w and $1 - p_w$, respectively, and $\text{Exp}(\tau)$ follows an exponential distribution with mean $1/\tau > 0$. Note that c can be viewed as a contamination parameter. The larger c is, the more contaminated the data. On the other hand, τ determines the asymmetry of the incidental parameters. The regression coefficients, $\beta_1^* = \dots = \beta_{50}^* = 1 \{ (X_{i,1}, \dots, X_{i,50}) \}^{i.i.d.} N(0, \Sigma_X)$, where $\Sigma_X(i, j) = 2 \exp(-|i - j|)$, which is a Toeplitz matrix and the constant 2 is used to inflate the covariance a little; the covariates are independent of $\{\epsilon_i\}^{i.i.d.} N(0, 1)$; and $n = 500$; $p_0 = 0.8$, $p_1 = 0.1$, $p_2 = 0.1$, c is 0.5, 1, 3 or 5, is 0.5 or 0.75 and $\tau = 1$.

5.1 Performance of Penalized Methods

The following methods for estimating β^* are evaluated. (i) Oracle method (O): an oracle knows the index set S of zero μ_i^* 's. Its performance is used as a benchmark. (ii) Ordinary least squares method (OLS): all μ_i^* 's are thought as zeros. (iii) Four penalized least squares (PLS) methods, namely, PLS with soft penalty (PLS.Soft or S), PLS with hard penalty (PLS.Hard or H), two-step PLS with soft penalty (PLS.Soft.TwoStep or S.TS) and two-step PLS with hard penalty (PLS.Hard.TwoStep or H.TS). More specifically, the oracle estimator of μ^* is given by $\hat{\beta}^{(O)} = (\sum_{i \in S} X_i X_i^T)^{-1} \sum_{i \in S} X_i Y_i$. The hard penalty function is $p_\lambda(|t|) = \lambda^2 - (|t| - \lambda)^2 \{ |t| < \lambda \}$ (see Fan and Li (2001)). Each method is evaluated by the square root of the empirical mean squared error (RMSE). Every penalized method is evaluated with a grid of values for the regularization parameter λ , ranging from .5 to 5 by .25.

The sequence plot of Figure 1 shows 500 realized incidental parameters μ_i^* 's with $c = 3$ and $p_w = 0.75$. They are used in data generation of simulations. The scatter plot of Figure 1 shows the responses Y_i 's against the first covariate X_{i1} 's of a generated data set and the red stars stand for the contaminated sample points, the ones with nonzero μ_i^* 's. With fifty covariates, usually it is difficult to graphically identify the contaminated data points.

With the above incidental parameters, those six methods are evaluated by simulations with iteration number 1000. Since Σ_X is a Toeplitz matrix with equal diagonal elements, the

asymptotic variances of the estimators of β_1^* and β_2^* , are different and representative for estimators of other β_i^* . So, we only report simulation results on the estimation of, β_1^* and β_2^* .

Figure 2 shows RMSE's of six estimators for, β_1^* . RMSE's for, β_2^* are similar. As expected, the oracle method has the smallest RMSE and OLS the largest. RMSE of each PLS method as a function of λ forms a convex curve, which achieves a minimal RMSE significantly below the green line of OLS and close to the cyan line of O. More specifically, RMSE of PLS.Hard achieves the minimal RMSE when λ is around 2.75. On the other hand, RMSE of PLS.Soft decreases a little till λ is around .75, then increases and stays above RMSE of PLS.Hard. This reflects the fact that a large λ in a soft-threshold method usually causes bias. PLS.Hard.TwoStep has very similar performance with PLS.Hard for all λ .

PLS.Soft.TwoStep has similar performance with PLS.Soft when λ is small. However, as λ becomes large, PLS.Soft.TwoStep moves closer to PLS.Hard than PLS.Soft. It is because PLS.Soft.TwoStep and PLS.Hard.TwoStep have similar estimation when λ is large. The minimal RMSE of PLS.Soft is slightly larger than those of other PLS Methods.

Table 1 depicts the minimal RMSE's of the estimators for β_1^* and β_2^* with the corresponding optimal λ 's and biases. The biases are ignorable comparing with the RMSE's. The optimal λ for PLS.Soft and other PLS methods are around .75 and 2.5, respectively. This indicates the simple soft threshold method tends to work best with a small λ due to the bias issue. Denote the empirical relative efficiency (ERE) of an estimator λ with respect to another estimator B as $\text{RMSE}(B)/\text{RMSE}(A)$. Then, for the estimation of β_1^* , the ERE's of PLS.Soft, PLS.Hard, PLS.Soft.TwoStep and PLS.Hard.TwoStep with respect to O are around 78%, 87%, 87% and 87%, respectively; the ERE's of the PLS methods with respect to OLS are around 176%, 196%, 196% and 196%, respectively. The ERE's for β_2^* are similar. Thus, in terms of ERE (and RMSE), the PLS methods perform closely to O and significantly better than OLS.

From Table 1, we can also see that the RMSE's of the estimators for β_1^* are always smaller than those for β_2^* . This is because the first covariate is less correlated with others covariates than the second one.

In order to examine the performance of the methods with general incidental parameters but not just those in Figure 1, we generate μ^* randomly for each iteration. The iteration number for each simulation is also 1000.

Figure 3 shows the RMSE's of six estimators of β_1^* under two settings: $p_w = 0.5$ and $c = 1$ or 5. Each plot in Figure 3 presents a similar pattern with Figure 2. When p_w is fixed at 0.5, the RMSE's of each non-oracle estimator of β_1^* increases as the contamination parameter c increases from 1 to 5. This indicates that each non-oracle estimator performs worse as the data becomes more contaminated. However, the PLS estimators are more robust than OLS, which is very sensitive to the change of c . We have also done simulations with $p_w = 0.75$ and

the RMSE's of the estimators of are similar to those with $p_w = 0.75$ so that the corresponding plots are similar to those in Figure 3. In other words, the RMSE's of all estimators are stable with respect to p_w , which means the magnitudes of the nonzero incidental parameters matter most but not their signs. We also note that some penalized methods perform closely to or even outperform the oracle one when c is small as showed in the plot with $c = 1$. This happens because that O ignores all the contaminated data points, even those with very light contamination, but the penalized methods exploit information in such points.

Table 2 contains the RMSE's of β_1^* the estimators of under eight settings with $p_w = 0.5$ or 0.75 and $c = 0.5, 1, 3$ or 5 . For each p_w , as c increases from 0.5 to 5 , the RMSE's of O with the multiplication factor 10^2 is almost constantly around 4 , those of OLS increases from about 4 to 8.5 and those of PLS ones grow from about 4 to 5 , which confirms the robustness of the PLS estimators. When $c = 1$ is small with respect to the variance of random error $\sigma = 1$, the data points are only slightly contaminated. OLS and PLS methods perform similar to O. However, when $c = 3$ is large, which means the data are more contaminated, the RMSE's of OLS become significantly larger, but PLS methods perform still closely to O.

5.2 Performance of Data-Driven Penalized Methods

Previous simulations have shown the PLS methods with optimal λ 's have good RMSE's comparing with those of the Oracle and OLS ones. In practice, however, these optimal λ 's are unknown. One approach to obtain a data driven λ has been introduced in Subsection 3.2. Since, as shown in the previous simulation results, the two-step PLS methods perform similarly with the one-step PLS methods, i.e. PLS.Soft and PLS.Hard, only the latter are studied by simulations with data-driven λ 's and denoted as PLS.Soft.Prac (S.P) and PLS.Hard.Prac (H.P), respectively. For estimating data-driven λ 's, $\alpha_I = 2$ and $\alpha_U = 7$. The size of the pure data set n_{pure} is $n/2$ and that of the testing data set is $n_{pure}/2$.

Simulations are first run with the deterministic sparse incidental parameters as showed in Figure 1. We can see in Table 1 that the RMSE's of estimators of β_1^* from PLS.Soft.Prac and PLS.Hard.Prac are around 5.4 and 5.0 , slightly larger than the optimal values 5.1 and 4.6 , respectively. However, they are still significantly smaller than RMSE of OLS, which is 8.4 . The observations of the estimators of β_2^* are similar. As before, we also evaluate the performance of the data-driven PLS methods with random sparse incidental parameters. Table 2 shows that, for a given p_w , when c is small such as 0.5 and 1 , the RMSE's of PLS.Soft.Prac and PLS.Hard.Prac are close to those of PLS.Soft and PLS.Hard with the optimal λ 's. In these cases, PLS.Soft.Prac performs slightly better than PLS.Hard.Prac, and even better than PLS.Soft with the optimal λ and the Oracle method. On the other hand, for a given, when c is large such as 3 and 5 , the RMSE's of PLS.Soft.Prac and PLS.Hard.Prac are greater than those of PLS.Soft and PLS.Hard, respectively, but still less than those of OLS. In these cases, the RMSE's of PLS.Soft.Prac are larger than those of PLS.Hard.Prac, which indicates the bias issue of the soft threshold method. Thus, the data-driven regularization parameter works well with penalized estimation. When the data is slightly contaminated, the soft penalty is preferred; otherwise, the hard penalty is recommended.

Tables 1 and 2 also contain RMSE of the least absolute deviation regression method (LAD) used in Fan et al. (2012b) with *all* but not part of the sample points with small residuals. Generally speaking, in both deterministic and random incidental parameter cases, LAD performs similarly with the PLS methods with data-driven λ 's. More specifically, when c is small, PLS.Soft.Prac outperform LAD; otherwise, LAD performs better. For all the cases, LAD is dominated by PLS.Hard.Prac. These observations confirm that LAD is an effective robustness method and the penalized methods make improvement.

5.3 Data-Driven Confidence Intervals

We next turn to investigate the finite-sample performance of the asymptotic confidence interval (CI) (3.5) for β_j^* with $j = 1, 2$ based on PLS two-step methods. Since CI (3.5) is based on the properties of the penalized two-step estimator with the soft penalty, we focus on PLS.TS.Soft with a data-driven regularization parameter λ . The choice of λ in Subsection 3.2 for minimizing RMSE is usually no longer suitable for constructing confidence intervals, since it is designed to achieve minimal RMSE. We propose to first obtain $\hat{\sigma}_{pure}$ as in the data-driven procedure in Subsection 3.2 and then simply set the data-driven λ be five times of $\hat{\sigma}_{pure}$. Since $\hat{\sigma}_{pure}$ tends to underestimate σ , this data-driven λ is usually not large with respect to σ . Denote this method as PLS.TwoStage.Soft.Prac or S.TS.P. After plugging in $\hat{\sigma}$ and $\hat{\sigma}_j^{-1}$, the square root of the (j, j) th element of $\hat{\Sigma}_X^{-1}$, and replace n by $m = \#(\hat{I}_0)$ in the theoretical CI (3.5), we obtain a data-driven CI $[\tilde{\beta}_j \pm m^{-1/2} \hat{\sigma} \hat{\sigma}_j^{-1} z_{\alpha/2}]$, where $\tilde{\beta}_j$ is the PLS.TwoStage.Soft.Prac estimator of β_j^* for each j .

This data-driven CI is compared with CI's based on Oracle and OLS methods. More specifically, denote the Oracle and OLS estimators of β_j^* as $\hat{\beta}_j^{(O)}$ and $\hat{\beta}_j^{(OLS)}$, respectively. Then, the corresponding CI's are given by $[\hat{\beta}_j^{(O)} \pm m_o^{-1/2} \hat{\sigma}^{(O)} \hat{\sigma}_j^{-1} z_{\alpha/2}]$ and $[\hat{\beta}_j^{(OLS)} \pm n^{-1/2} \hat{\sigma}^{(OLS)} \hat{\sigma}_j^{-1} z_{\alpha/2}]$, where m_o is the number of zero incidental parameters and $\hat{\sigma}^{(O)}$ and $\hat{\sigma}^{(OLS)}$ are the estimators of σ from O and OLS methods, respectively.

The simulation settings are the same to the previous ones with *deterministic* sparse incidental parameters except the following changes. (a) The number of covariates d is reduced to 5 from 50. This is because when $d = 50$ and the nominal level is 95%, even the empirical coverage rate (CR) of the oracle confidence interval for β_1^* becomes 93.5%, not very close to 95%. (b) The iteration number is increased from 1000 to 10000 to improve the accuracy of CR's. (c) The probabilities of nonzero incidental parameters (p_1, p_2) are set to be (0.01, 0.01), (0.03, 0.03) and (0.05, 0.05); the contamination parameter c is increased to 10. In order to achieve good second order asymptotic approximation, we can either increase the sample size or enlarge the signal noise ratio. Here we adopt the latter.

Table 3 reports the empirical coverage rates (CR) and average lengths (AL) of the CI's of β_1^* and β_2^* from O, OLS and PLS.TS.Soft.Prac methods under three different settings on the

incidental parameters. For the oracle method, these three settings are the same and thus only one set of simulation results are presented. Table 3 shows that the CR's of all methods under all settings are close to the nominal level .95. The OLS treats the deterministic incidental parameters as random ones and achieves excellent CR's. However, the AL's of OLS are significantly larger than those of O and PLS.TS.Soft.Prac, especially when there are more non-zero incidental parameters. On the other hand, the AL's of PLS.TS.Soft.Prac are only slightly larger than those of O. This means PLS.TS.Soft.Prac has excellent efficiency in terms of AL's given excellent CR's. Also note that the AL's for β_1^* are less than those for β_2^* . This is because the asymptotic variance of $\hat{\beta}_1$ is less than that of $\hat{\beta}_2$ when the covariance matrix X is a Toeplitz matrix. Simulations with random incidental parameters under the same settings have also been done and the results are similar to those in Table 3 with slightly inflated AL's for OLS and PLS.TS.Soft.Prac due to the randomness of the incidental parameters.

5.4 Real Data Analysis

We implement the penalized estimation with the soft penalty in the method of estimating false discovery proportion of a multiple testing procedure proposed by Fan et al. (2012b) for investigating the association between the expression level of gene CCT8, which is closely related to Down Syndrome phenotypes, and thousands of SNPs. The data set consists of three populations: 60 Utah residents (CEU), 45 Japanese and 45 Chinese (JPTCHB) and 60 Yoruba (YRI). More details on the data set can be found in Fan et al. (2012b).

In the testing procedure by Fan et al. (2012b), a filtered least absolute deviation regression (LAD) is used to estimate the loading factors with 90% of the cases (SNPs) whose test statistics are small and thus the resulting estimator is statistically biased. We upgrade this step with S.P described in Subsections 3.2 and 5.2 and re-estimate the number of false discoveries $V(t)$ and the false discovery proportion FDP(t) as functions of $-\log_{10}(t)$, where t is a thresholding value. Figure 4 shows the number of total discoveries $R(t)$, $\hat{V}(t)$ and $\widehat{\text{FDP}}(t)$ from procedures using filtered LAD and S.P. It is clear that $\hat{V}(t)$ and $\widehat{\text{FDP}}(t)$ with S.P are uniformly larger than but reasonably close to those with filtered LAD. Table 4 contains $R(t)$ and $\widehat{\text{FDP}}(t)$ with filtered LAD and S.P for several specific thresholds. The estimated FDPs with S.P for CEU and YRI are slightly larger than those with LAD and $\widehat{\text{FDP}}$ for JPTCHB with S.P is more than double of that with filtered LAD. This suggests that the estimation of FDP with filtered LAD might tend to be optimistic.

6 Conclusion and Discussion

This paper considers the estimation of structural parameters with a finite or diverging number of covariates in a linear regression model with the presence of high-dimensional sparse incidental parameters. By exploiting the sparsity, we propose an estimation method penalizing the incidental parameters. The penalized estimator of the structural parameters is consistent and asymptotically Gaussian and achieves an oracle property. On the contrary, the penalized estimator of the incidental parameters possesses only partial selection consistency but not consistency. Thus, the structural parameters are consistently estimated while the

incidental parameters not, which presents a partial consistency phenomenon. Further, in order to construct better confidence regions for the structural parameters, we propose a two-step estimator, which has fewer possible asymptotic distributions and can be asymptotically even more efficient than the one-step penalized estimator when the size and magnitude of nonzero incidental parameters are substantially large.

Simulation results show that the penalized methods with best regularization parameters achieve significantly smaller mean square errors than the ordinary least squares method which ignores the incidental parameters. Also provided is a data-driven regularization parameter, with which the penalized estimators continue to significantly outperform ordinary least squares when the incidental parameters are too large to be neglected. In terms of average length together with excellent coverage rates, the advantage of the confidence intervals based on the two-step estimator with an alternative data-driven regularization parameter is verified by simulations. A data set on genome-wide association is analyzed with a multiple testing procedure equipped with a data-driven penalized method and false discovery proportions are estimated.

In econometrics, a fixed effect panel data model is given by, for $1 \leq i \leq n$ and $1 \leq t \leq T$

$$Y_{it} = \mu_i^* + X_{it}^T \beta^* + \epsilon_{it}, \quad (6.1)$$

where μ_i^* 's are unknown fixed effects. When T diverges, the fixed effects can be consistently estimated. When T is finite and greater than or equal to 2, although the fixed effects can no longer be consistently estimated, they can be removed by a within-group transformation: for each i , $Y_{it} - \bar{Y}_i = (X_{it} - \bar{X}_i)^T \beta^* + \epsilon_{it} - \bar{\epsilon}_i$ where $\bar{Y}_i$, $\bar{X}_i$ and $\bar{\epsilon}_i$ are the averages of Y_{it} 's, X_{it} 's and ϵ_{it} 's, respectively. When T is equal to 1, however, the within-group transformation fails. Note that, with $T = 1$, Model (6.1) becomes Model (1.1) so that the proposed penalized estimations provide a solution under the sparsity assumption on the fixed effects.

Although this paper only illustrates the partial consistency phenomenon of a penalized estimation method for a linear regression model, such a phenomenon shall universally exist for a general parametric model, which contains both a structural parameter and a high-dimensional sparse incidental parameter. For example, consider a panel data logistic regression model: $P(Y_{it} = 1 | X_{it}) = (1 + \exp\{-(\mu_i^* + X_{it}^T \beta^*)\})^{-1}$. When T is finite, the fixed effects μ_i^* 's cannot be removed by the within-group transformation as in the panel data linear model (6.1). However, the proposed penalized estimations can still provide a solution.

Further, if the structural parameter has a dimension diverging faster than the sample size and is sparse, it is expected that the partial consistency phenomenon will continue to appear when sparsity penalty is imposed on both the structural and incidental parameters.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgement

The authors thank the Editor, an associate editor, and three referees for their many helpful comments that have resulted in significant improvements in the article. The research was partially supported by NSF grants DMS-1206464 and DMS-0704337 and NIH Grants R01-GM100474-01 and R01-GM072611-08.

A: Appendix

In this appendix, we provide the proofs of the theoretical results in Section 4. The proofs of the results in Sections 2 and 3 are in Supplements D and E.

Denote $\mathbb{S}_{k,l} = \mathbb{S}_{\{k,k+1,\dots,l\}}$ and $\mathbb{S}_{k,l}^c = \mathbb{S}_{\{k,k+1,\dots,l\}}^c$. Let $\mathcal{B} = \left\{ \max_{s+1 \leq i \leq n} \|X_i\|_2 \leq \kappa_n \right\}$ and $\mathcal{D} = \bigcap_{i=1}^n \{-\gamma_n \leq \epsilon_i \leq \gamma_n\}$. Then $P(\mathcal{B}) \rightarrow 1$ and $P(\mathcal{D}) \rightarrow 1 \rightarrow$ by (2.2)

Proof of Lemma 4.1. We first consider S_{j0} 's, then S_{j1} 's, and finally S_{j2} 's with $j = 1, 2, 3$.

Consider S_{10} , S_{20} and S_{30} . Let $\mathcal{A} = \{S_{10} = S_{10}^*\}$. Note that $P(\mathcal{A}) \geq P(\mathcal{A}|\mathcal{B})P(\mathcal{B})$ and

$P(\mathcal{B}) \rightarrow 1$. It suffices to show that $P(\mathcal{A}|\mathcal{B}) \rightarrow 1$. By $\lambda \gg \sqrt{d}\kappa_n$, it follows

$$P(\mathcal{A}|\mathcal{B}) \geq P\left(\left\{s+1 \leq i \leq n: -\lambda + \max_{s+1 \leq i \leq n} \|X_i\|_2 \sqrt{d}C \leq \epsilon_i \leq \lambda - \max_{s+1 \leq i \leq n} \|X_i\|_2 \sqrt{d}C\right\} \cap S_{10}^*|\mathcal{B}\right) \geq P\left(\left\{s+1 \leq i \leq n: -\lambda + \kappa_n \sqrt{d}C \leq \epsilon_i \leq \lambda - \kappa_n \sqrt{d}C\right\} \cap S_{10}^*\right) \geq P(\mathcal{D}) \rightarrow 1$$

Thus, wpg1, $S_{10} = S_{10}^*$. From $S_{10} \cup S_{20} \cup S_{30} = S_{10}^*$, it follows that, wpg1, $S_{20} = S_{30} = \emptyset$.

Consider S_{21} , S_{31} and S_{11} . Recall that $\mu^* = \min\{\mu_i^*: 1 \leq i \leq s_1\}$ and note that

$\lambda - \mu^* + \sqrt{d}C\kappa_n < -\gamma_n$ when n is large. Let $S_{211} = S_{21} \cap S_{21}^*$ and $S_{212} = S_{21} \cap S_{21}^{*c}$. We will

show $P(S_{211} = S_{21}^*) \rightarrow 1$ and $P(S_{212} = \emptyset) \rightarrow 1$. Then $P(S_{21} = S_{21}^*) \rightarrow 1$. Denote

$\mathcal{A}_1 = \{S_{211} \supset S_{21}^*\}$. On the event $\mathcal{B}$, $S_{211} \supset \{1 \leq i \leq s_1: \epsilon_i > \lambda - \mu^* + \sqrt{d}C\kappa_n \text{ and}$

$\mu_i^* > 0\} \supset \{1 \leq i \leq s_1: \epsilon_i > -\gamma_n \text{ and } \mu_i^* > 0\}$. Then,

$$P(\mathcal{A}_1) \geq P(\mathcal{A}_1|\mathcal{B})P(\mathcal{B}) \geq P(\{1 \leq i \leq s_1: \epsilon_i > -\gamma_n \text{ and } \mu_i^* > 0\} \cap S_{21}^*)P(\mathcal{B}) \rightarrow 1 \cdot 1 = 1. \text{ It}$$

follows that, wpg1, $S_{211} \supset S_{21}^*$. Note that $S_{211} \subset S_{21}^*$. Then, wpg1, $S_{211} = S_{21}^*$. Denote

$\mathcal{A}_2 = \{S_{212} = \emptyset\}$. On the event $\mathcal{B}$,

$S_{12} \supset \{s_1 + 1 \leq i \leq s: -\lambda - \mu_i^* + \sqrt{d}C\kappa_n \leq \epsilon_i \leq \lambda - \mu_i^* - \sqrt{d}C\kappa_n\}$, which contains

$\{s_1 + 1 \leq i \leq s: -\gamma_n \leq \epsilon_i \leq \gamma_n\}$. Then

$$P(\mathcal{A}) \geq P(\mathcal{A}|\mathcal{B})P(\mathcal{B}) \geq P(\{s_1 + 1 \leq i \leq s: -\gamma_n \leq \epsilon_i \leq \gamma_n\} = S_{12}^*)P(\mathcal{B}) \rightarrow 1. \text{ Then, wpg1,}$$

$P(S_{21} = S_{21}^*) \rightarrow 1$. Similarly, we can show, wpg1, $S_{31} = S_{31}^*$. Note that S_{11} , S_{21} and S_{31} are

disjoint and their union is $S_{21}^* \cup S_{31}^*$. Then, wpg1, $S_{11} = \emptyset$. Consider S_{12} , S_{22} and S_{32} .

Denote $\mathcal{A} = \{S_{12} = S_{12}^*\}$. Note that $-\lambda - \mu_i^* + \sqrt{d}C\kappa_n < -\gamma_n$ and $\lambda - \mu_i^* - \sqrt{d}C\kappa_n > \gamma_n$ when n

is large for $s_1 + 1 \leq i \leq s$. On the event $\mathcal{B}$,

$S_{12} \supset \{s_1 + 1 \leq i \leq s: -\lambda - \mu_i^* + \sqrt{d}C\kappa_n \leq \epsilon_i \leq \lambda - \mu_i^* - \sqrt{d}C\kappa_n\}$, which contains

$\{s_1 + 1 \leq i \leq s: -\gamma_n \leq \epsilon_i \leq \gamma_n\}$. Then,

$P(\mathcal{A}) \geq P(\mathcal{A}|\mathcal{B})P(\mathcal{B}) \geq P(\{s_1 + 1 \leq i \leq s: -\gamma_n \leq \epsilon_i \leq \gamma_n\} = S_{12}^*)P(\mathcal{B}) \rightarrow 1$. Thus, wpg1, $S_{12} = S_{12}^*$. Note that S_{12} , S_{22} and S_{32} are disjoint and their union is S^*_{12} . Then, wpg1, $S_{22} = S_{32} = \emptyset$. $\square$

Before proceeding to the proofs of Theorems 4.2 to A.6, we denote $\bar{\sigma}_X^2 = (1/d) \sum_{j=1}^d \text{Var}[X_{0j}]$ and $\bar{\sigma}_{XX}^2 = (1/d^2) \sum_{k=1}^d \sum_{l=1}^d \text{Var}[X_{0k}X_{0l}]$ and make the following assumptions.

Assumption (E1): $\bar{\sigma}_X^2$ is bounded.

Assumption (E2): $\bar{\sigma}_{XX}^2$ is bounded.

Assumption (E) in Section 4 implies Assumptions (E1) and (E2) by Cauchy-Schwarz inequality. For simplicity, we adopt the notation $\lesssim$, which means the left hand side is bounded by a constant times the right, where the constant does not affect related analysis. Below are three lemmas needed for proving Theorems 4.2 to A.6. Their proofs are in

Supplement F. Suppose that $\mathbf{M}$ and $\mathbf{E}$ are matrices and $\|\cdot\|$ is a matrix norm and that $\{\mathbf{A}_n\}$ is

a sequence of random $d \times d$ matrices and $\mathbf{A}$ a deterministic $d \times d$ matrix, and denote $\hat{\Sigma}_n = (1/n)\mathbb{S}_n$, the sample covariance matrix.

Lemma A.1 (Stewart (1969)). If $\|\mathbf{I}\| = 1$ and $\|\mathbf{M}^{-1}\|\|\mathbf{E}\| < 1$, then

$$\frac{\|(\mathbf{M} + \mathbf{E})^{-1} - \mathbf{M}^{-1}\|}{\|\mathbf{M}^{-1}\|} \leq \frac{\|\mathbf{M}^{-1}\|\|\mathbf{E}\|}{1 - \|\mathbf{M}^{-1}\|\|\mathbf{E}\|}.$$

Lemma A.2. If $\|\mathbf{A}^{-1}\|_{F,d}$ is bounded, $\mathbf{A}_n \xrightarrow{P} \mathbf{A}$, and $r_d \geq 1/\sqrt{d}$, then $\mathbf{A}_n^{-1} \xrightarrow{P} \mathbf{A}^{-1}$, where the convergence in probability is wrt $rd\|\cdot\|_F$.

Lemma A.3. If Assumption (E2) holds and $r_d^2 d^4/n \rightarrow 0$, then $\hat{\Sigma}_n \xrightarrow{P} \Sigma_X$ wrt $rd\|\cdot\|_F$.

Proof of Theorem 4.2. By the proof of Lemma 4.1, wpg1, the solution $\hat{\beta}_n$ to $\varphi_n(\beta) = 0$ on

$\mathcal{B}_C(\beta^*)$ is explicitly given by $\hat{\beta}_n = \beta^* + T_0^{-1}(T_1 + T_2 + T_3 - T_4)$, where

$$T_0 = (1/n)\mathbb{S}_{s_1+1,n}, T_1 = (1/n)\mathbb{S}_{S_{12}^*}^\mu, T_2 = (1/n)\mathbb{S}_{s_1+1,n}^\epsilon, T_3 = (\lambda/n)\mathcal{S}_{S_{21}^*} \text{ and } T_4 = (\lambda/n)\mathcal{S}_{S_{31}^*}.$$

Then, $r_d\|\hat{\beta}_n - \beta^*\|_2 \leq \|T_0^{-1}\|_{F,d} \sum_{i=1}^4 r_d \sqrt{d} \|T_i\|_2$. We will show that $\|T_0^{-1}\|_{F,d}$ is bounded by a positive constant wpg1 and $r_d \sqrt{d} \|T_i\|_2 \xrightarrow{P} 0$ for $i = 1, 2, 3, 4$. Then, $r_d\|\hat{\beta}_n - \beta^*\|_2 = o_P(1)$.

Consider T_0 . By Lemma A.3, $\|T_0 - \Sigma_X\|_{F,d} \xrightarrow{P} 0$ under Assumption (E2) and the condition $d^3/n \rightarrow 0$. Then, by Lemma A.2, together with Assumption (D), $\|T_0^{-1} - \Sigma_X^{-1}\|_{F,d} \xrightarrow{P} 0$. This implies that, wpg1, $\|T_0^{-1}\|_{F,d}$ is bounded by a positive constant. Consider T_1 . Wpg1, $r_d \sqrt{d} \|T_1\|_2 \leq r_d \sqrt{d} s_2 \kappa_n \gamma_n / n = o(1)$ for $s_2 = o(n/(r_d \sqrt{d} \kappa_n \gamma_n))$. Consider T_2 . For any $\delta > 0$, $P(\|T_2\|_2 > \delta) \leq (1/\delta^2) P\left\| (1/n) \sum_{i=s_1+1}^n X_i \epsilon_i \right\|_2^2 \leq d \sigma^2 \bar{\sigma}_X^2 / (n \delta^2)$, where $\bar{\sigma}_X^2 = (1/d) \sum_{j=1}^d \sigma_j^2$. Thus, $P(r_d \sqrt{d} \|T_2\|_2 > \delta) \leq r_d^2 d^2 \sigma^2 \bar{\sigma}_X^2 / (n \delta^2) \rightarrow 0$ by Assumption (E1) and $(r_d d)^2/n \rightarrow 0$. Consider T_3 and T_4 . Wpg1, $r_d \sqrt{d} \|T_3\|_2 \leq r_d \sqrt{d} s_1 \kappa_n / n = o(1)$ for $s_1 = o(n/(r_d \sqrt{d} \lambda \kappa_n))$. Similarly $r_d \sqrt{d} \|T_4\|_2 = o_P(1)$. $\square$

The next lemma is needed for proving Theorem 4.3 and its proof is in Supplement F. Suppose $\{\xi_i\}$ are i.i.d. copies of ξ_0 , a d -dimensional random vector with mean zero. Denote $\sigma_{\xi, \max}^2 = \max_{1 \leq j \leq d} \text{Var}[\xi_{0j}]$, $\sigma_{\xi, \min}^2 = \min_{1 \leq j \leq d} \text{Var}[\xi_{0j}]$ and $\gamma_{\xi, \max} = \max_{1 \leq j \leq d} \mathbb{E}|\xi_{0j}|^3$.

Lemma A.4. Suppose $\sigma_{\xi, \max}$ and $\gamma_{\xi, \max}$ are bounded from above and $\sigma_{\xi, \max}$ is bounded from zero. If $d = o(\sqrt{n})$, then $(1/\sqrt{n}) \sum_{i=1}^n \xi_i = O_P(\sqrt{d \log d})$ wrt $\|\cdot\|_2$.

Proof of Theorem 4.3. We reuse the notations T_i 's in the proof of Theorems 4.2, from which, $\sqrt{n} A_n (\hat{\beta}_n - \beta^*) = V_1 + V_2 + V_3 - V_4$, where $V_i = B_n T_i$ for $i = 1, 2, 3, 4$ and $B_n = \sqrt{n} A_n T_0^{-1}$. It is sufficient to show that $V_2 \xrightarrow{d} N(0, \sigma^2 G_X)$ and other V_i 's are $o_P(1)$. Consider V_1 . We have

$\|V_1\|_2 \leq \sqrt{nd} \|A_n\|_F \|T_0^{-1}\|_{F,d} \|T_1\|_2$. By Assumption (F), $\|A_n\|_F$ is bounded. By Lemmas A.2 and A.3 and Assumption (D), for $d = o(n^{1/3})$, wpg1, $\|T_0^{-1}\|_{F,d}$ is bounded. We have, wpg1, $\|T_1\|_2 \leq s_2 \kappa_n \gamma_n / n$. Then, $\|V_1\|_2 \lesssim \sqrt{d} n s_2 \kappa_n \gamma_n$. Thus, $\|V_1\|_2 = o_P(1)$ for $s_2 = o(\sqrt{n}/(\sqrt{d} \kappa_n \gamma_n))$.

Consider V_2 . We have $V_2 = V_{21} + V_{22}$, where

$V_{21} = \sqrt{n} A_n \Sigma_X^{-1} T_2$ and $V_{22} = \sqrt{n} A_n (T_0^{-1} - \Sigma_X^{-1}) T_2$. First, note that

$V_{21} = \sqrt{(n-s_1)/n} \sum_{i=s_1+1}^n Z_{n,i}$, where $Z_{n,i} = (1/\sqrt{n-s_1}) A_n \Sigma_X^{-1} X_i \epsilon_i$. On one hand, for every

$\delta > 0$, $\sum_{i=s_1+1}^n \mathbb{E} \|Z_{n,i}\|_2^2 \mathbb{I}(\|Z_{n,i}\|_2 > \delta) \leq (n-s_1) \mathbb{E} \|Z_{n,0}\|_2^4 / \delta^2$, and

$$\mathbb{E} \|Z_{n,0}\|_2^4 = \frac{1}{(n-s_1)^2} \mathbb{E} \epsilon_0^4 \mathbb{E} (X_0^T \Sigma_X^{-1} A_n^T A_n \Sigma_X^{-1} X_0)^2, \text{ which is } \leq \frac{d^2}{(n-s_1)^2} \mathbb{E} \epsilon_0^4 \lambda_{\max}(G_n) \lambda_{\min}^{-2}(\Sigma_X) \kappa_X^2.$$

Then, by Assumptions (D'), (E) and (F) and for $d = o(\sqrt{n})$,

$\sum_{i=s_1+1}^n \mathbb{E} \|Z_{n,i}\|_2^2 \mathbb{I}(\|Z_{n,i}\|_2 > \delta) \rightarrow 0$. On the other hand

$\sum_{i=s_1+1}^n \text{Cov}(Z_{n,i}) = \sigma^2 A_n \Sigma_X^{-1} A_n^T \rightarrow \sigma^2 G_X$ by Assumption (F). Thus, by central limit

theorem (see Proposition 2.27 in van der Vaart (1998)), $V_{21} \xrightarrow{d} N(0, \sigma^2 G_X)$. Next, consider

V_{22} . Note that $\|V_{22}\|_2 \leq \|A_n\|_F (d \log(d))^{1/2} \|T_0^{-1} - \Sigma_X^{-1}\|_F (d \log(d))^{-1/2} \|\sqrt{n}T_2\|_2$. By Assumption (F), $\|A_n\|_F$; by Lemmas A.2 and A.3 $(d \log(d))^{1/2} \|T_0^{-1} - \Sigma_X^{-1}\|_F$ is $o_P(1)$ for $d^5 \log(d) = o(n)$; by Lemma A.4, $(d \log(d))^{-1/2} \|\sqrt{n}T_2\|_2 = (d \log(d))^{-1/2} \left\| \frac{1}{\sqrt{n}} \mathbb{S}_{s_1+1,n}^\epsilon \right\|_2$ is $O_P(1)$ for $d = o(\sqrt{n})$. Then, $V_{22} \xrightarrow{P} 0$. Thus, by Slutsky's lemma, $V_2 \xrightarrow{d} N(0, \sigma^2 G_X)$. Consider V_3 and V_4 . First consider V_3 . By noting that $s_1 = o(\sqrt{n}l(\lambda\sqrt{d}\kappa_n))$, wpgl,

$$\|V_3\|_2 \leq \sqrt{nd} \|A_n\|_F \|T_0^{-1}\|_{F,d} \|T_3\|_2 \lesssim \sqrt{d} \lambda s_1 \kappa_n / \sqrt{n} \rightarrow 0. \text{ Thus, } \|V_3\|_2 = o_P(1). \text{ In the same way, } \|V_4\|_2 = o_P(1). \square$$

Proof of Theorem 4.4. By the definition of ϵ , we have $P(\epsilon) = T_1 T_2 T_3$, where

$$T_1 = P\left(\cap_{i=1}^{s_1} \left\{ \left| \mu_i^* + X_i^T(\beta^* - \hat{\beta}) + \epsilon_i \right| > \lambda \right\}\right), T_2 = P\left(\cap_{i=s_1+1}^s \left\{ \left| \mu_i^* + X_i^T(\beta^* - \hat{\beta}) + \epsilon_i \right| \leq \lambda \right\}\right)$$

and $T_3 = P\left(\cap_{i=s+1}^n \left\{ \left| X_i^T(\beta^* - \hat{\beta}) + \epsilon_i \right| \leq \lambda \right\}\right)$. We will show that each T_i converges to one.

Then, $P(\mathcal{E}) \rightarrow 1$. Denote $c = \{r_d \|\hat{\beta} - \beta^*\|_2 \leq 1\}$. Then $P(\mathcal{E}) \rightarrow 1$ since 3 is a consistent estimator of β^* wrt $r_d \|\cdot\|_2$. Consider T_1 . We have $1 - T_1 \leq T_{11} + T_{12}$, where

$$T_{11} = P\left(\cup_{i \in S_{21}^*} \left\{ \left| \mu_i^* + X_i^T(\beta^* - \hat{\beta}) + \epsilon_i \right| \leq \lambda \right\}\right) \text{ and}$$

$$T_{12} = P\left(\cup_{i \in S_{31}} \left\{ \left| \mu_i^* + X_i^T(\beta^* - \hat{\beta}) + \epsilon_i \right| \leq \lambda \right\}\right). \text{ It is sufficient to show that both } T_{11} \text{ and } T_{12}$$

converge to zero. By

$$\sqrt{d}\kappa_n \ll \lambda \ll \mu^*, T_{11} \leq P\left(\cup_{i \in S_{21}^*} \left\{ \epsilon_i \leq \lambda - \mu^* + \|X_i\|_2 \cdot \|\hat{\beta} - \beta^*\|_2 \right\}, \mathcal{E}\right) + P(\mathcal{E}^c), \text{ which is}$$

$$\leq P\left(\cup_{i \in S_{21}^*} \left\{ \epsilon_i \leq \lambda - \mu^* + \sqrt{d}\kappa_n \right\}\right) + P(\mathcal{E}^c) \leq s_1 P\{\epsilon_0 \leq -\gamma_n\} + P(\mathcal{E}^c) \rightarrow 0. \text{ Similarly,}$$

$T_{12} \rightarrow 0$. Thus $T_1 \rightarrow 1$. Consider T_2 and T_3 . By $\alpha\gamma_n \leq \lambda$ and

$$\sqrt{d}\kappa_n \ll \lambda, T_2 \geq P\left(\cap_{i=s_1}^s \left\{ -\lambda - \mu_i^* + (1/r_d)\kappa_n \leq \epsilon_i \leq \lambda - \mu_i^* - (1/r_d)\kappa_n \right\}, \mathcal{E}\right), \text{ which is}$$

$$\geq P\left(\cap_{i=s_1}^s \left\{ -\lambda - \mu_i^* + \sqrt{d}\kappa_n \leq \epsilon_i \leq \lambda - \mu_i^* - \sqrt{d}\kappa_n \right\}, \mathcal{E}\right) \geq P\left(\cap_{i=s_1}^s \left\{ -\gamma_n \leq \epsilon_i \leq \gamma_n \right\}, \mathcal{E}\right) \rightarrow 1.$$

Then $T_2 \rightarrow 1$. Similarly, $T_3 \rightarrow 1$. $\square$

Proof of Theorem 4.5. Note that $\sqrt{n}A_n \Sigma_X^{1/2}(\tilde{\beta} - \beta^*) = \tilde{R}_1 + \tilde{R}_2 + V_1 + V_2$, where $\tilde{R}_1 = \sqrt{n}A_n R_1$,

$$\tilde{R}_2 = \sqrt{n}A_n R_2, R_1 = \left(X_{\hat{I}_0}^T X_{\hat{I}_0}\right)^{-1} X_{\hat{I}_0}^T Y_{\hat{I}_0} \{\hat{I}_0 \neq I_0\}, R_2 = -\left(X_{I_0}^T X_{I_0}\right)^{-1} X_{I_0}^T Y_{I_0} \{\hat{I}_0 \neq I_0\}, \text{ and } V_i \text{'s}$$

are defined in the proof of Theorem 4.3. Since $P(\|\tilde{R}_1\|_2 = 0) \geq P\{\hat{I}_0 = I_0\} \rightarrow 1$, we have

$\tilde{R}_1 = o_P(1)$. Similarly, $\tilde{R}_2 = o_P(1)$. By the proof of Theorem 4.3,

$V_1 = o_P(1)$ and $V_2 \xrightarrow{d} N(0, \sigma^2 G_X)$. Therefore, the desired result follows by Slutsky's lemma.

$\square$

Lemma A.5 (Consistency on $\hat{\sigma}$). *Suppose the assumptions and conditions of Theorem 4.2 hold with $r_d \geq \sqrt{d}$. If $s_2 = o(n/\gamma_n^2)$, then $\hat{\sigma} \xrightarrow{P} \sigma$.*

Proof of Lemma A.5. Since the assumptions and conditions of Theorem 4.2 hold with $r_d \geq \sqrt{d}$, the penalized estimators $\hat{\beta}$ and $\tilde{\beta}$ are consistent estimators of β^* wrt $\sqrt{d} \|\cdot\|_2$ by Theorems 4.2 and F.4 in Supplement F. Let $\mathcal{A} = \{\hat{I}_0 = I_0\}$. Then $\mathcal{A}$ occurs wpg1 by Theorem 4.4.

Note that $\hat{\sigma}^2 = T\mathcal{A} + \hat{\sigma}^2\mathcal{A}^c$, where $T = (n - s_1)^{-1} \|Y_{I_0} - X_{I_0}^T \tilde{\beta}\|_2^2$. It suffices to show that

$$T \xrightarrow{P} \sigma^2. \text{ Note that } T = \sum_{i=1}^6 T_i, \text{ where } T_1 = (n - s_1)^{-1} \sum_{i=s_1+1}^n [X_i^T (\beta^* - \tilde{\beta})]^2,$$

$$T_2 = (n - s_1)^{-1} \sum_{i=s_1+1}^n \epsilon_i^2, T_3 = 2(n - s_1)^{-1} \sum_{i=s_1+1}^n X_i^T (\beta^* - \tilde{\beta}) \epsilon_i,$$

$$T_4 = (n - s_1)^{-1} \sum_{i=s_1+1}^s \mu_i^{*2}, T_5 = 2(n - s_1)^{-1} \sum_{i=s_1+1}^s \mu_i X_i^T (\beta^* - \tilde{\beta}) \text{ and}$$

$$T_6 = 2(n - s_1)^{-1} \sum_{i=s_1+1}^s \mu_i^* \epsilon_i. \text{ It is clear that } T_2 \xrightarrow{P} \sigma^2. \text{ Thus, it is sufficient to show other}$$

T_i 's are $o_P(1)$. For every $\eta > 0$, wpg1, $\sqrt{d} \|\beta^* - \tilde{\beta}\|_2 \leq \eta$. By Assumption (E1), wpg1,

$$|T_1| \leq \frac{1}{d} \frac{1}{n - s_1} \sum_{i=s_1+1}^n \|X_i^T\|_2^2 (\sqrt{d} \|\beta^* - \tilde{\beta}\|_2)^2 \leq 2\eta^2 \frac{1}{d} \mathbb{E} \|X_0^T\|_2^2 = 2\eta^2 \bar{\sigma}_X^2 \lesssim \eta^2. \text{ For every}$$

$$\eta > 0, \text{ wpg1, } |T_3| \leq 2 \frac{1}{\sqrt{d} n - s_1} \sum_{i=s_1+1}^n \|X_i^T \epsilon_i\|_2 \sqrt{d} \|\beta^* - \tilde{\beta}\|_2 \leq 4\eta \frac{1}{\sqrt{d}} \mathbb{E} \|X_0^T \epsilon_0\|_2 = 4\sigma\eta \bar{\sigma}_X \lesssim \eta.$$

For $s_2 = o(n/\gamma_n^2)$, $|T_4| \leq (n - s_1)^{-1} s_2 \gamma_n^2 \rightarrow 0$. For $s_2 = o(\sqrt{dn}/(\gamma_n \kappa_n))$, wpg1,

$$|T_5| \leq 2 \frac{1}{\sqrt{d} n - s_1} s_2 \gamma_n \kappa_n \sqrt{d} \|\beta^* - \tilde{\beta}\|_2 \leq 2\eta \frac{1}{\sqrt{d} n - s_1} s_2 \gamma_n \kappa_n \xrightarrow{P} 0. \text{ For } s_2 = o(n/\gamma_n), \text{ wpg1, } |T_6|$$

$$\leq 4 \frac{1}{n - s_1} \gamma_n s_2 \mathbb{E} |\epsilon_0| \rightarrow 0.$$

Theorem A.6 (Asymptotic Distributions on $\hat{\beta}$ and $\hat{\beta}$ with $\hat{G}_{X,n}$). *Under the assumptions and conditions of Theorem 4.3, if $d^8(\log(d))^2 = o(n)$, then $\sqrt{n} \hat{G}_{X,n}^{-1/2} A_n(\hat{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 I_q)$.*

Similarly, under the assumptions and conditions of Theorem 4–5, If $d^8(\log(d))^2 = o(n)$, then

$$\sqrt{n} \hat{G}_{X,n}^{-1/2} A_n(\tilde{\beta} - \beta^*) \xrightarrow{d} N(0, \sigma^2 I_q).$$

Note that a stronger requirement on d is required to handle $\hat{G}_{X,n}^{-1/2}$ in Theorem A.6.. Below is a lemma needed for proving Theorem A.6.

Lemma A.7 (Wahler (2009)). *Suppose A and B are $m \times m$ symmetric positive-semidefinite matrices. Then, for $p > 1$, $\|A^{1/p} - B^{1/p}\|_F^p \leq m^{(p-1)/2} \|A - B\|_F^p$. Specifically, for*

$$p = 2, \|A^{1/2} - B^{1/2}\|_F \leq (m^{1/2} \|A - B\|_F)^{1/2}.$$

Proof of Theorem A.6. We only show the result on $\hat{\beta}$. since the result on β . can be obtained in a similar way. We reuse the notations T_i 's in the proof of Theorems 4.2, from which, $\sqrt{n}\hat{G}_{X,n}^{-1/2}A_n(\hat{\beta}_n - \beta^*) = M + R$ where $M = \sqrt{n}G_{X,n}^{-1/2}A_n(\hat{\beta}_n - \beta^*)$ and $R = \sqrt{n}(\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2})A_n(\hat{\beta}_n - \beta^*)$. By Theorem 4.3, $M \xrightarrow{d} N(0, \sigma^2 G_X)$. Then, it is sufficient to show that $R \xrightarrow{P} 0$ wrt $\|\cdot\|_2$. We have $R = R_1 + R_2 + R_3 - R_4$, where $R_i = B_n T_i$ for $i = 1, 2, 3, 4$ and $B_n = \sqrt{n}(\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2})A_n T_0^{-1}$. We will show each R_i converges to zero in probability, which finishes the proof. Before that, we first establish an inequality for $\|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F$. By Lemma A.7, $\|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F \leq (\sqrt{q}\|\hat{G}_{X,n}^{-1} - G_{X,n}^{-1}\|_F)^{1/2}$. Note that, by Lemma A.3 $\|\hat{\Sigma}_n - \Sigma_X\|_F \xrightarrow{P} 0$ for $d^4 = o(n)$. Then, by Lemma A.2, $\|\hat{G}_{X,n} - G_{X,n}\|_F \leq \|A_n\|_F^2 \|\hat{\Sigma}_n^{-1} - \Sigma_X^{-1}\|_F \lesssim \|A_n\|_F^2 \|\hat{\Sigma}_n - \Sigma_X\|_F \xrightarrow{P} 0$. Thus, by Lemma A.2, $\|\hat{G}_{X,n}^{-1} - G_{X,n}^{-1}\|_F \lesssim \|\hat{G}_{X,n} - G_{X,n}\|_F \lesssim \|A_n\|_F^2 \|\hat{\Sigma}_n - \Sigma_X\|_F$. Since q is a fixed integer, it follows $\|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F \lesssim \|A_n\|_F (\sqrt{q}\|\hat{\Sigma}_n - \Sigma_X\|_F)^{1/2} \lesssim \|A_n\|_F (\|\hat{\Sigma}_n - \Sigma_X\|_F)^{1/2}$. Consider R_1 . Note that $\|R_1\|_2 \leq \sqrt{n}\sqrt{d}\|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F \|A_n\|_F \|T_0^{-1}\|_{F,d} \|T_1\|_2$, which is $\lesssim \sqrt{n}(d\|\hat{\Sigma}_n - \Sigma_X\|_F)^{1/2} \|A_n\|_F^2 \|T_0^{-1}\|_{F,d} \|T_1\|_2$. By Lemmas A.2 and A.3, $d\|\hat{\Sigma}_n - \Sigma_X\|_F = OP(1)$ for $d^6 = o(n)$. By Assumption (F), $\|A_n\|_F$ is $O(1)$. is bounded. By Lemmas A.2 and A.3 and Assumption (D), for $d = o(n^{1/3})$, $\|T_0^{-1}\|_{F,d}$ is bounded. Also note that, $\|T_1\|_2 \leq s_2 \kappa_n \gamma_n / n$. Then, $\|R_1\|_2 \lesssim s_2 \kappa_n \gamma_n / \sqrt{n}$. Thus, $\|R_1\|_2 = o_P(1)$ for $s_2 = o(\sqrt{n}/(\kappa_n \gamma_n))$. Consider R_2 . Note that $\|R_2\|_2 \leq \|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F \|A_n\|_F \|T_0^{-1}\|_{F,d} \|\sqrt{n}T_2\|_2$, which is $\lesssim (d^2 \log(d)\|\hat{\Sigma}_n - \Sigma_X\|_F)^{1/2} \|A_n\|_F^2 \|T_0^{-1}\|_{F,d} (\log(d))^{-1/2} \|\sqrt{n}T_2\|_2$. By Lemmas A.2 and A.3, $d^2 \log(d)\|\hat{\Sigma}_n - \Sigma_X\|_F$ is $OP(1)$ for $d^8 (\log(d))^2 = o(n)$. By Assumption (F), $\|A_n\|_F$ is $O(1)$. By Lemmas A.2 and A.3 and Assumption (D), for $d = o(n^{1/3})$, $\|T_0^{-1}\|_{F,d}$ is bounded. By Lemma A.4, $(d \log(d))^{-1/2} \|\sqrt{n}T_2\|_2 = (d \log(d))^{-1/2} \left\| \frac{1}{\sqrt{n}} \mathbb{S}_{s_1+1,n}^e \right\|_2$ is $O_P(1)$ is bounded. By Lemma A.4, $(d \log(d))^{-1/2} \|\sqrt{n}T_2\|_2 = (d \log(d))^{-1/2} \left\| \frac{1}{\sqrt{n}} \mathbb{S}_{s_1+1,n}^e \right\|_2$ for $d = o(\sqrt{n})$. Thus, $R_2 \xrightarrow{P} 0$. Consider R_3 and R_4 . By $s_1 = o(\sqrt{n}/(\lambda \kappa_n))$, $\|R_3\|_2 \leq \sqrt{n} \|\hat{G}_{X,n}^{-1/2} - G_{X,n}^{-1/2}\|_F \|A_n\|_F \|T_0^{-1}\|_{F,d} \|T_3\|_2$, which is $\lesssim \sqrt{n}(d\|\hat{\Sigma}_n - \Sigma_X\|_F)^{1/2} \|A_n\|_F^2 \|T_0^{-1}\|_{F,d} \|T_3\|_2 \lesssim \lambda s_1 \kappa_n / \sqrt{n} \rightarrow 0$. Thus, $\|R_3\|_2 = o_P(1)$. In the same way, $\|R_4\|_2 = o_P(1)$.

References

- Basu Debabrata. On the elimination of nuisance parameters. *Journal of the American Statistical Association*, 72:355–366, 1977.
- Bean Derek, Bickel Peter, Noureddine El Karoui Chingway Lim, and Yu Bin. Penalized robust regression in high-dimension. 2012.
- Louis HY Chen Larry Goldstein, and Man Shao. *Normal Approximation by Stein’s Method*. Springer-Verlag, 2010.
- Fan Jianqing and Li Runze. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 12 2001.
- Fan Jianqing and Lv Jinchi. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20:101–148, 2010. [PubMed: 21572976]
- Fan Jianqing and Peng Heng. On non-concave penalized likelihood with diverging number of parameters. *The Annals of Statistics*, 32:928–961, 2004.
- Fan Jianqing, Liao Yuan, and Mincheva Martina. High dimensional covariance matrix estimation in approximate factor models. 2011.
- Fan Jianqing, Fan Yingying, and Barut Emre. Adaptive robust variable selection. 2012a.
- Fan Jianqing, Feng Yang, and Tong Xin. A road to classification in high dimensional space: the regularized optimal affine discriminant. *Journal of the Royal Statistical Society Series B*, 2012b.
- Huber Peter J. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35:73–101, 1964.
- Huber Peter J. Robust regression: Asymptotics, conjectures and monte carlo. *The Annals of Statistics*, 1:799–821, 1973.
- Jahn Johannes. *Introduction to the Theory of Nonlinear Optimization*. Springer Berlin Heidelberg, 2007.
- Kiefer J and Wolfowitz J Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics*, 27:887–906, 1956.
- Kosorok Michael R. Bootstrapping the grenander estimator In *Beyond Parametrics in Interdisciplinary Research: Festschrift in Honor of Professor Sen Pranab K.*, pages 282–292. Institute of Mathematical Statistics: Hayward, CA., 2008.
- Lambert-Lacroix S and Zwald L Robust regression through the huber’s criterion and adaptive lasso penalty. *Electronic Journal of Statistics*, 5:1015–1053, 2011.
- Lancaster Tony. The incidental parameter problem since 1948. *Journal of Econometrics*, 95: 391–413, 2000.
- Moreira Marcelo J. A maximum likelihood method for the incidental parameter problem. *The Annals of Statistics*, 37(6A):3660–3696, 2009.
- Neyman Jerzy and Scott Elizabeth L. Consistent estimates based on partially consistent observations. *Econometrica*, 16:1–32, 1948.
- Portnoy Stephen and He Xuming. A robust journey in the new millennium. *Journal of the American Statistical Association*, 95:1331–1335, 2000.
- Shiryayev Albert N. *Probability*. Springer-Verlag, second edition, 1995.
- Stewart GW On the continuity of the generalized inverse. *SIAM Journal on Applied Mathematics*, 17:33–45, 1969.
- Tang Runlong, Banerjee Moulinath, and Kosorok Michael R. Likelihood based inference for current status data on a grid: A boundary phenomenon and an adaptive inference procedure. *The Annals of Statistics*, 40(1):45–72, 2012.
- Tibshirani Robert. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc B*, 58: 267–288, 1996.
- Aad W van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- Aad W van der Vaart and Wellner Jon A. *Weak Convergence and Empirical Processes*. Springer, 1996.
- Wihler Thomas P. On the holder continuity of matrix functions for normal matrices. *Journal of inequalities in pure and applied mathematics*, 10, 2009.

- Zhao Peng and Yu Bin. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7(Nov):2541–2563, 2006.
- Fan Jianqing and Lv Jinchi. Non-concave penalized likelihood with np-dimensionality. *IEEE Transactions On Information Theory*, 57:5467–5484, 2011. [PubMed: 22287795]
- Fan Jianqing and Peng Heng. On non-concave penalized likelihood with diverging number of parameters. *The Annals of Statistics*, 32:928–961, 2004.
- Kosorok Michael R. *Introduction to Empirical Processes and Semiparametric Inference*. Springer New York, 2008.
- Neyman Jerzy and Scott Elizabeth L. Consistent estimates based on partially consistent observations. *Econometrica*, 16:1–32, 1948.
- Shiryaev Albert N. *Probability*. Springer-Verlag, second edition, 1995.
- Aad W van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- Zhao Peng and Yu Bin. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7(Nov):2541–2563, 2006.

Figure 1:

The sequence plot on the left shows 500 incidental parameters's with $c = 3$ and $p_w = 0.75$.

The non-zero' μ_i^* s are in red. The scatter plot on the right shows the responses Y_j 's against the first covariate X_{1j} 's of a data set generated with those 500 incidental parameters. The red stars stand for the contaminated sample points, the ones with nonzero μ_i^* 's.

Figure 2: RMSE's of the O(Oracle), OLS, S(PLS.Soft), H(PLS.Hard), S.TH(PLS.Soft.TwoStep) and H.TH(PLS.Hard.TwoStep) estimators of β_1^* with the incidental parameters shown in Figure 1. The top and bottom solid horizontal lines show the RMSE's for OLS and O, respectively. Other four horizontal lines indicate the minimal RMSE's for those four PLS methods and the corresponding four best λ 's are shown by the vertical lines.

Figure 3:

Similar to Figure 2, these plots show the RMSE's of the O(Oracle), OLS, S(PLS.Soft), H(PLS.Hard), S.TH(PLS.Soft.TwoStep) and H.TH(PLS.Hard.TwoStep) estimators of β_1^* with randomly generated μ^* under two settings with $p_w = 0.5$ and $c = 1$ or 5 .

Figure 4: Discovery number $R(t)$ estimated false discovery number $V(t)$ and estimated false discovery proportion $FDP(t)$ as functions of a threshold on t for populations CEU, JPTCHB and YRI. The x -axis is $-\log_{10}(t)$.

Table 1:

RMSE's of the Oracle, OLS and LAD estimators and The minimal RMSE's of the penalized estimators of β_1^* and β_2^* with the corresponding optimal or data-driven λ 's and biases when the incidental parameters shown in Figure 1 are used. For a data-driven method, the data-driven λ is different in each iteration so that the reported λ 's are averages. The lines over the numbers emphasize the numbers are averages. The standard deviations for the data-driven λ 's of S.P and H.P are 0.37 and 0.34, respectively.

	O	OLS	S	H	S.TS	H.TS	S.P	H.P	LAD
Bias($\hat{\beta}_1$) $\times 10^4$	.8	17.8	1.3	.6	5.9	20.4	9.5	11.5	10.4
RMSE($\hat{\beta}_1$) $\times 10^2$	4.0	8.4	5.1	4.6	4.6	4.6	5.4	5.0	5.4
λ			.75	2.75	2.25	2.5	<u>2.45</u>	<u>2.47</u>	
Bias($\hat{\beta}_2$) $\times 10^4$	-2.3	-46.7	24.6	43.5	-21.8	10.0	32.3	13.9	-7.3
RMSE($\hat{\beta}_2$) $\times 10^2$	4.4	9.0	5.3	4.9	5.0	4.9	5.8	5.5	5.9
λ			.75	2.75	2.75	2.5	<u>2.45</u>	<u>2.45</u>	

Table 2:

Similar to Table 1, this one shows the RMSE's and minimal RMSE's of nine estimators of β_1^* and β_2^* under eight settings on randomly generated μ^* with $p_w = 0.5, 0.75$ and $c = 0.5, 1, 3, 5$. The standard deviations of the data-driven λ 's of S.P and H.P for different settings are between 0.2 and 0.45.

$\text{RMSE}(\hat{\beta}_1) \times 10^2$	O	OLS	S	H	S.TS	H.TS	S.P	H.P	LAD
$(p_w, c) = (.5, .5)$	4.06	4.06	3.92	3.95	3.90	3.96	4.01	4.38	4.85
λ			1.25	3.5	2.75	2.5	$\overline{2.08}$	$\overline{2.15}$	
$(p_w, c) = (.5, 1)$	4.06	4.58	4.11	4.27	4.21	4.27	4.01	4.59	4.75
λ			2	3.75	4.25	3.5	$\overline{2.20}$	$\overline{2.24}$	
$(p_w, c) = (.5, 3)$	4.12	6.47	4.81	4.99	4.81	4.80	5.64	5.19	5.49
λ			1	2.5	2	2.25	$\overline{2.42}$	$\overline{2.45}$	
$(p_w, c) = (.5, 5)$	4.07	8.50	4.96	4.63	4.66	4.64	6.36	4.87	5.52
λ			1	2.25	2.75	3	$\overline{2.52}$	$\overline{2.58}$	
$(p_w, c) = (.75, .5)$	4.13	4.02	3.91	3.88	3.98	3.91	4.08	4.41	4.79
λ			1.5	5	3.25	4.5	$\overline{2.08}$	$\overline{2.14}$	
$(p_w, c) = (.75, 1)$	4.17	4.41	4.19	4.15	4.15	4.20	4.16	4.59	4.98
λ			3.25	3	4	3.5	$\overline{2.15}$	$\overline{2.23}$	
$(p_w, c) = (.75, 3)$	4.15	5.99	4.91	4.93	4.80	5.02	5.35	5.06	5.52
λ			1	2.25	2	2.5	$\overline{2.44}$	$\overline{2.47}$	
$(p_w, c) = (.75, 5)$	3.97	8.41	5.01	4.66	4.75	4.66	6.22	4.90	5.76
λ			1.5	2.25	2.5	3	$\overline{2.55}$	$\overline{2.60}$	

Table 3:

Coverage rates (CR) and average length (AL) of 95% confidence intervals for β_1^* and β_2^* from O, OLS, PLS.TwoStage.Soft.Prac methods under three settings on deterministic sparse incidental parameters.

(p_1, p_2)	β_1^*	O	OLS	S.TS.P	β_2^*	O	OLS	S.TS.P
(.01,.01)			.950	.944			.948	.945
(.03,.03)	CR	.955	.951	.948	CR	.948	.946	.947
(.05,.05)			.946	.945			.949	.946
(.01,.01)			.200	.135			.213	.144
(.03,.03)	AL	.133	.279	.137	AL	.142	.297	.146
(.05,.05)			.382	.139			.407	.149

Table 4:

Discovery numbers $R(t)$ and estimated false discover proportions $\widehat{\text{FDP}}(t)$ from methods with LAD and S.P for specific values of threshold t .

Population	t	$R(t)$	$\widehat{\text{FDP}}(t)$ with LAD	$\widehat{\text{FDP}}(t)$ with S.P
CEU	6.12×10^{-4}	4	.810	.845
JPTCHB	1.51×10^{-9}	5	.153	.373
YRI	2.54×10^{-9}	2	.227	.308