

Model-Free Feature Screening and FDR Control With Knockoff Features

Wanjun Liu^a, Yuan Ke^b, Jingyuan Liu^c, and Runze Li^d

^aDepartment of Statistics, The Pennsylvania State University, University Park, PA; ^bDepartment of Statistics, University of Georgia, Athens, GA; ^cMOE Key Laboratory of Econometrics, Department of Statistics, School of Economics, Wang Yanan Institute for Studies in Economics, and Fujian Key Lab of Statistics, Xiamen University, Xiamen, China; ^dDepartment of Statistics, The Pennsylvania State University, University Park, PA

ABSTRACT

This article proposes a model-free and data-adaptive feature screening method for ultrahigh-dimensional data. The proposed method is based on the projection correlation which measures the dependence between two random vectors. This projection correlation based method does not require specifying a regression model, and applies to data in the presence of heavy tails and multivariate responses. It enjoys both sure screening and rank consistency properties under weak assumptions. A two-step approach, with the help of knockoff features, is advocated to specify the threshold for feature screening such that the false discovery rate (FDR) is controlled under a prespecified level. The proposed two-step approach enjoys both sure screening and FDR control simultaneously if the prespecified FDR level is greater or equal to $1/s$, where s is the number of active features. The superior empirical performance of the proposed method is illustrated by simulation examples and real data applications. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received August 2019

Accepted June 2020

KEYWORDS

Data adaptive; Multivariate response model; Nonlinear model; Projection correlation; Rank consistency; Sure screening

1. Introduction

The technological development has made extensive data collection and storage feasible in diverse fields. Datasets with ultrahigh-dimensional features characterize many contemporary research problems in machine learning, computer science, statistics, engineering, social science, finance, and so on. When the features contain redundant or noisy information, estimating their functional relationship with the response may become quite challenging in terms of computational expediency, statistical accuracy, and algorithmic stability (Fan, Samworth, and Wu 2009; Hall and Miller 2009; Lv and Liu 2014). To overcome such challenges caused by ultrahigh-dimensionality, Fan and Lv (2008) proposed a sure independence screening (SIS) method, which aims to screen out the redundant features by ranking their marginal Pearson correlations. The SIS method is named after the SIS property, which states the selected subset of features contains all the active ones with probability approaching one. The promising numerical performance soon made SIS popular among ultrahigh-dimensional studies (Liu, Zhong, and Li 2015; Liu and Li 2020). The sure screening idea has been applied to many important statistical problems including generalized linear model (Fan and Song 2010), multi-index semiparametric model (Zhu et al. 2011), nonparametric model (Fan, Feng, and Song 2011; Liu, Li, and Wu 2014), quantile regression (He, Wang, and Hong 2013; Wu and Yin 2015), and compressed sensing (Xue and Zou 2011), among others.

In addition to the sure screening property, we argue an appealing screening method should satisfy the following two properties. First, the screening method should be model-free in the sense that it can be implemented without specifying a regression model. In the ultrahigh-dimensional regime, it is challenging, if not impossible, to specify a correct regression model before removing a huge number of redundant features. Hence, the model-free property is desired as it guarantees the effectiveness of the screening method in the presence of model misspecification. The model-free screening method becomes a hot research topic in recent years (Zhu et al. 2011; Li, Zhong, and Zhu 2012; Mai and Zou 2015). The second property is data-adaptive which means the screening method should not be sensitive to assumptions like independence, sub-Gaussianity, and univariate response. Such assumptions are usually not satisfied under ultrahigh-dimensional settings. Even met on the population level, they can be violated in the observed sample due to ultrahigh-dimensionality. Therefore, the screening methods that are sensitive to such assumptions may perform poorly in real applications. The data-adaptive screening methods also draw a certain amount of attention recently. For instance, He, Wang, and Hong (2013), Wu and Yin (2015), and Ma, Li, and Tsai (2017), among others, considered quantile-based screening which adapts to heavy-tailed data. In addition, Wang (2012) and Fan, Ke, and Wang (2020) developed screening methods for strongly correlated features.

CONTACT Yuan Ke  yuan.ke@uga.edu  Department of Statistics, University of Georgia, Athens, GA 30602; Jingyuan Liu  jingyuan@xmu.edu.cn  MOE Key Laboratory of Econometrics, Department of Statistics, School of Economics, Wang Yanan Institute for Studies in Economics, and Fujian Key Lab of Statistics, Xiamen University, Xiamen 361005, China.

All authors equally contributed to this work, and the authors are listed in seniority.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JASA.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

© 2020 American Statistical Association

Unfortunately, none of the aforementioned screening methods enjoys sure screening, model-free, and data-adaptive properties simultaneously. For example, the SIS is tailored to the linear regression and depends on the independence assumption. Li, Zhong, and Zhu (2012) developed a model-free SIS procedure based on the distance correlation. However, its sure screening property requires sub-exponential assumption for features and response. The Kolmogorov distance based screening method proposed in Mai and Zou (2012) is robust against heavy-tailed data, but it only works for binary classification problems. Pan, Wang, and Li (2016) proposed a pairwise sure screening procedure for linear discriminant analysis, which is not sensitive to the tail behavior of the features, but requires balanced categories.

In this article, we propose a model-free and data-adaptive feature screening method named PC-Screen. The PC-Screen is based on ranking the projection correlations between features and response variables. The projection correlation, proposed by Zhu et al. (2017), is a measure of dependence between two random vectors which enjoys several nice probability properties. The PC-Screen does not require specifying any regression model and is insensitive to the correlations and moment conditions of the dataset. As the projection correlation is dimension-free to both random vectors, the PC-Screen can be applied to multitask learning problems (Caruana 1997). For instance, we can find a parsimonious set of features that are jointly dependent on the multivariate response. As existing asymptotic results are not suitable for the “large p small n problem” studied in this article, we also develop nonasymptotic concentration-type inequalities for empirical projection correlation. Further, we demonstrate the proposed PC-Screen enjoys not only the sure screening property but also a stronger result called rank consistency property. The only condition required is a minimum signal strength gap between active and inactive features. The extensive numerical examples show the proposed method wins the horse racing against its competitors in various scenarios.

Most feature screening methods depend on some threshold parameter that controls the cut-off between active and inactive features. The optimal choice of the threshold typically depends on unknown parameters. Under a specific model assumption, the threshold can be selected by cross-validation or information criteria approaches. However, when no model assumption is imposed, such parameter selection approaches are not applicable as the loss function which measures the goodness of fit is not well defined. In addition, existing screening methods tend to sacrifice the false discovery rate (FDR) for sure screening property by choosing a conservative threshold parameter, leading to an inflated model size. Recently, Zhao and Li (2012) studied the SIS for Cox models with a principled method to select the threshold aiming at controlling the false positive rate. Song et al. (2014) proposed a censored rank independence screening for survival data and chose the threshold by estimating the proportion of active features. The validity of this procedure relies on the independence assumption of multiple test statistics which may easily get violated in the model-free setup. In this article, we tackle the issue of threshold selection with a two-step procedure named PC-Knockoff. In the first step, we apply the PC-Screen method to obtain an over-fitted subset of

moderate size from the ultrahigh-dimensional features. In the second step, we construct knockoff counterparts for the features which survive in the first step. Conditioning on the sure screening of the first step, we further select a parsimonious model with FDR controlled below a prespecified level with a statistic that uses the knockoff features. Theoretical analysis shows that when the prespecified level is not too aggressive, the PC-Knockoff procedure enjoys sure screening property, as well as conditional FDR control simultaneously with high probability. We also validate the theoretical findings with various numerical examples.

The rest of the article is organized as follows. In Section 2, we briefly review the definition and properties of projection correlation and demonstrate its nonasymptotic properties. Then, we propose the PC-Screen procedure and show its sure screening and rank consistency properties under very mild conditions. Section 3 studies the PC-Knockoff procedure which selects the threshold with FDR control as well as sure screening property. Section 4 assesses the finite sample performance of the PC-Screen and PC-Knockoff methods with several simulated examples and a real data application. We briefly summarize the article in Section 5. Due to the limited space, we provide the proofs of some theoretical results in Appendix A, and relegate the proofs of remaining results as well as some additional numerical examples to the supplementary materials.

2. Model-Free and Data-Adaptive Screening Procedure

2.1. Projection Correlation

To pave the way for the proposed screening procedure, we first provide some background on the projection correlation and its properties introduced in Zhu et al. (2017). Let $\mathbf{x} \in \mathbb{R}^p$ and $\mathbf{y} \in \mathbb{R}^q$ be two random vectors. The projection correlation is elicited by the following independence testing problem,

$$H_0 : \mathbf{x} \text{ and } \mathbf{y} \text{ are independent} \quad \text{versus} \quad H_1 : \text{otherwise.}$$

The null hypothesis holds if and only if $U = \boldsymbol{\alpha}^T \mathbf{x}$ and $V = \boldsymbol{\beta}^T \mathbf{y}$ are independent for all unit vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. Let $F_{U,V}(u, v)$ be the joint distribution of (U, V) , and $F_U(u)$ and $F_V(v)$ be the marginal distributions of U and V . The squared projection covariance is defined as

$$\begin{aligned} \text{Pcov}(\mathbf{x}, \mathbf{y})^2 &= \iiint (F_{U,V}(u, v) - F_U(u)F_V(v))^2 dF_{U,V}(u, v) d\boldsymbol{\alpha} d\boldsymbol{\beta} \\ &= \iiint \text{cov}^2\{I(\boldsymbol{\alpha}^T \mathbf{x} \leq u), I(\boldsymbol{\beta}^T \mathbf{y} \leq v)\} dF_{U,V}(u, v) d\boldsymbol{\alpha} d\boldsymbol{\beta}, \end{aligned} \quad (1)$$

where $I(\cdot)$ is the indicator function. Furthermore, the projection correlation between \mathbf{x} and \mathbf{y} is defined as the square root of

$$\text{PC}(\mathbf{x}, \mathbf{y})^2 = \frac{\text{Pcov}(\mathbf{x}, \mathbf{y})^2}{\text{Pcov}(\mathbf{x}, \mathbf{x})\text{Pcov}(\mathbf{y}, \mathbf{y})}, \quad (2)$$

and we follow the convention $0/0 = 0$.

In general $0 \leq \text{PC}(\mathbf{x}, \mathbf{y}) \leq 1$, testing whether \mathbf{x} and \mathbf{y} are independent amounts to testing whether $\text{PC}(\mathbf{x}, \mathbf{y}) = 0$. The projection correlation is a measure of dependence between two

random vectors and enjoys some appealing properties. Let \mathbf{x} and \mathbf{y} be two random vectors with continuous marginal and joint probability distributions, $\text{PC}(\mathbf{x}, \mathbf{y}) = 0$ if and only if \mathbf{x} and \mathbf{y} are independent. We remark here this property does not hold in general without the assumption that (\mathbf{x}, \mathbf{y}) is jointly continuous. When \mathbf{x} and \mathbf{y} are two dependent discrete random variables that are constructed in a similar fashion as in Hoeffding (1948), it is possible that we have $\text{PC}(\mathbf{x}, \mathbf{y}) = 0$.

Zhu et al. (2017) gave an explicit formula for the squared projection covariance in (1). Let $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_5, \mathbf{y}_5)$ be five independent random copies of (\mathbf{x}, \mathbf{y}) , then

$$\begin{aligned} & \text{Pcov}(\mathbf{x}, \mathbf{y})^2 \\ &= S_1 + S_2 - 2S_3 \\ &= E \left[\arccos \left\{ \frac{(\mathbf{x}_1 - \mathbf{x}_3)^T (\mathbf{x}_4 - \mathbf{x}_3)}{\|\mathbf{x}_1 - \mathbf{x}_3\| \|\mathbf{x}_4 - \mathbf{x}_3\|} \right\} \arccos \left\{ \frac{(\mathbf{y}_1 - \mathbf{y}_3)^T (\mathbf{y}_4 - \mathbf{y}_3)}{\|\mathbf{y}_1 - \mathbf{y}_3\| \|\mathbf{y}_4 - \mathbf{y}_3\|} \right\} \right] \\ &+ E \left[\arccos \left\{ \frac{(\mathbf{x}_1 - \mathbf{x}_3)^T (\mathbf{x}_4 - \mathbf{x}_3)}{\|\mathbf{x}_1 - \mathbf{x}_3\| \|\mathbf{x}_4 - \mathbf{x}_3\|} \right\} \arccos \left\{ \frac{(\mathbf{y}_2 - \mathbf{y}_3)^T (\mathbf{y}_5 - \mathbf{y}_3)}{\|\mathbf{y}_2 - \mathbf{y}_3\| \|\mathbf{y}_5 - \mathbf{y}_3\|} \right\} \right] \\ &- 2E \left[\arccos \left\{ \frac{(\mathbf{x}_1 - \mathbf{x}_3)^T (\mathbf{x}_4 - \mathbf{x}_3)}{\|\mathbf{x}_1 - \mathbf{x}_3\| \|\mathbf{x}_4 - \mathbf{x}_3\|} \right\} \arccos \left\{ \frac{(\mathbf{y}_2 - \mathbf{y}_3)^T (\mathbf{y}_4 - \mathbf{y}_3)}{\|\mathbf{y}_2 - \mathbf{y}_3\| \|\mathbf{y}_4 - \mathbf{y}_3\|} \right\} \right], \end{aligned} \quad (3)$$

where $\|\cdot\|$ is the L_2 norm. Equation (3) shows that the projection covariance only depends on the vectors through forms $(\mathbf{x}_k - \mathbf{x}_l)/\|\mathbf{x}_k - \mathbf{x}_l\|$ and $(\mathbf{y}_k - \mathbf{y}_l)/\|\mathbf{y}_k - \mathbf{y}_l\|$ whose second moments are unity. This gives us the intuition that the projection covariance is free of the moment conditions on (\mathbf{x}, \mathbf{y}) which are usually required by some other measurements, such as distance correlation (Li, Zhong, and Zhu 2012).

Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$ and $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^T$ be an observed sample of (\mathbf{x}, \mathbf{y}) . Equation (3) leads to a straightforward estimator of $\text{Pcov}(\mathbf{x}, \mathbf{y})^2$ based on a U -statistic, yet it is difficult to calculate (Székely and Rizzo 2010). An equivalent form of the U -statistic is given in Zhu et al. (2017). In particular, the squared sample projection variance and covariance between \mathbf{X} and \mathbf{Y} can be calculated as

$$\begin{aligned} \widehat{\text{Pcov}}(\mathbf{X}, \mathbf{Y})^2 &= n^{-3} \sum_{k,l,r=1}^n A_{klr} B_{klr}, \\ \widehat{\text{Pcov}}(\mathbf{X}, \mathbf{X})^2 &= n^{-3} \sum_{k,l,r=1}^n A_{klr}^2, \text{ and} \\ \widehat{\text{Pcov}}(\mathbf{Y}, \mathbf{Y})^2 &= n^{-3} \sum_{k,l,r=1}^n B_{klr}^2, \end{aligned} \quad (4)$$

where $k, l, r = 1, \dots, n$,

$$\begin{aligned} a_{klr} &= \arccos \left\{ \frac{(\mathbf{x}_k - \mathbf{x}_r)^T (\mathbf{x}_l - \mathbf{x}_r)}{\|\mathbf{x}_k - \mathbf{x}_r\| \|\mathbf{x}_l - \mathbf{x}_r\|} \right\}, \\ a_{klr} &= 0 \text{ if } k = r \text{ or } l = r, \\ \bar{a}_{k \cdot r} &= n^{-1} \sum_{l=1}^n a_{klr}, \\ \bar{a}_{\cdot l r} &= n^{-1} \sum_{k=1}^n a_{klr}, \\ \bar{a}_{\cdot \cdot r} &= n^{-2} \sum_{k=1}^n \sum_{l=1}^n a_{klr}, \end{aligned}$$

$$\begin{aligned} A_{klr} &= a_{klr} - \bar{a}_{k \cdot r} - \bar{a}_{\cdot l r} + \bar{a}_{\cdot \cdot r}, \\ b_{klr} &= \arccos \left\{ \frac{(\mathbf{y}_k - \mathbf{y}_r)^T (\mathbf{y}_l - \mathbf{y}_r)}{\|\mathbf{y}_k - \mathbf{y}_r\| \|\mathbf{y}_l - \mathbf{y}_r\|} \right\}, \\ b_{klr} &= 0 \text{ if } k = r \text{ or } l = r, \\ \bar{b}_{k \cdot r} &= n^{-1} \sum_{l=1}^n b_{klr}, \\ \bar{b}_{\cdot l r} &= n^{-1} \sum_{k=1}^n b_{klr}, \\ \bar{b}_{\cdot \cdot r} &= n^{-2} \sum_{k=1}^n \sum_{l=1}^n b_{klr}, \\ B_{klr} &= b_{klr} - \bar{b}_{k \cdot r} - \bar{b}_{\cdot l r} + \bar{b}_{\cdot \cdot r}. \end{aligned}$$

Then the sample projection correlation between \mathbf{X} and \mathbf{Y} is defined as the square root of

$$\widehat{\text{PC}}(\mathbf{X}, \mathbf{Y})^2 = \frac{\widehat{\text{Pcov}}(\mathbf{X}, \mathbf{Y})^2}{\widehat{\text{Pcov}}(\mathbf{X}, \mathbf{X}) \widehat{\text{Pcov}}(\mathbf{Y}, \mathbf{Y})}. \quad (5)$$

Based on (4), the sample projection correlation can be computed in $O(n^3)$.

We first provide exponential-type deviation inequalities for squared sample projection covariance and correlation.

Theorem 1. For any $0 < \varepsilon < 1$ satisfying $n \geq 10\pi^2/\varepsilon$, there exists positive constants c_1 and c_2 , such that

$$\Pr \{ |\widehat{\text{Pcov}}(\mathbf{X}, \mathbf{Y})^2 - \text{Pcov}(\mathbf{x}, \mathbf{y})^2| > \varepsilon \} \leq c_1 \exp\{-c_2 n \varepsilon^2\}, \text{ and}$$

$$\Pr \{ |\widehat{\text{PC}}(\mathbf{X}, \mathbf{Y})^2 - \text{PC}(\mathbf{x}, \mathbf{y})^2| > \varepsilon \} \leq 5c_1 \exp\{-c_2 \sigma n \varepsilon^2\},$$

where $\sigma = \min\{\sigma_x^3 \sigma_y^3 / 4M^4, \sigma_x^2 \sigma_y^2 / 4M^4\}$, $\sigma_x = \text{Pcov}(\mathbf{x}, \mathbf{x})^2$, $\sigma_y = \text{Pcov}(\mathbf{y}, \mathbf{y})^2$, and $M = 4\pi^2$.

The proof of Theorem 1 is based on an exponential-type deviation inequality for U -statistic and can be found in the supplementary materials. The above exponential inequalities do not depend on the dimensionality and moment conditions of both random vectors. The exception probability decays exponentially with sample size n which guarantees good finite sample performance of the proposed estimator.

2.2. PC-Screen Procedure

In this subsection, we propose a model-free and data-adaptive screening procedure using the nice properties of projection correlation. Let $\mathbf{y} = (Y_1, \dots, Y_q)^T$ be the response vector of q variables and $\mathbf{x} = (X_1, \dots, X_p)^T$ be the vector of p features. To avoid the trivial discussion, we restrict ourselves to the non-degenerate case. That is, $\min_{1 \leq k \leq p} \text{Pcov}(X_k, X_k)^2 \geq \sigma_0^2$ and $\min_{1 \leq k \leq q} \text{Pcov}(Y_k, Y_k)^2 \geq \sigma_0^2$ for some $\sigma_0 > 0$.

Denote $F(\mathbf{y}|\mathbf{x})$ the conditional distribution function of \mathbf{y} given \mathbf{x} . Without specifying any regression model of \mathbf{y} on \mathbf{x} , we define the index set of active features by

$$\mathcal{A} = \{k : F(\mathbf{y}|\mathbf{x}) \text{ functionally depends on } X_k, k = 1, \dots, p\}.$$

The number of active features is $s = |\mathcal{A}|$, the cardinality of \mathcal{A} . The features that do not belong to \mathcal{A} are inactive features.

We use \mathcal{A}^c , the complement of \mathcal{A} , to denote the index set of inactive features. The above setting abstracts a large number of sparse regression problems including linear model, generalized linear model, additive model, semiparametric model, nonlinear model, and so on. Moreover, it also allows the multivariate response and grouped predictors.

Suppose that $\{(\mathbf{x}_i, \mathbf{y}_i), i = 1, \dots, n\}$ is a random sample of (\mathbf{x}, \mathbf{y}) . Denote $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$ and $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^T$. In the ultrahigh-dimensional regime, it is natural to assume that the number of features p greatly exceeds the sample size n , but the number of active features s is smaller than n . For a given feature $X_k, k = 1, \dots, p$, a sufficient condition for X_k to be an inactive feature is the independence between X_k and \mathbf{y} . This, together with [Theorem 1](#), motivates us to screen out the features whose projection correlations with \mathbf{y} are close to 0. As a result, we estimate the set of active features by

$$\widehat{\mathcal{A}}(\delta) = \{k : \widehat{PC}(\mathbf{X}_k, \mathbf{Y})^2 \geq \delta, 1 \leq k \leq p\},$$

where δ is a prespecified positive threshold, and \mathbf{X}_k is the k th column of \mathbf{X} . With a proper choice of δ , we show that the proposed feature screening procedure enjoys the sure screening property, which states that with probability approaching 1, all active features are included in $\widehat{\mathcal{A}}(\delta)$. We name this feature screening procedure as projection correlation based screening, or PC-Screen.

Denote $\omega_k \equiv PC(X_k, \mathbf{y})^2$ and $\widehat{\omega}_k \equiv \widehat{PC}(\mathbf{X}_k, \mathbf{Y})^2$ to be the squared population and sample projection correlations between the k th feature and response, respectively. To analyze the property of PC-Screen, we impose the following minimum signal strength condition.

Condition 1 (Minimum signal strength).

- (a) For some $c_3 > 0$ and $0 \leq \kappa < 1/2$, $\min_{k \in \mathcal{A}} \omega_k \geq 2c_3 n^{-\kappa}$.
- (b) For some $c_3 > 0$ and $0 \leq \kappa < 1/2$, $\min_{k \in \mathcal{A}} \omega_k - \max_{k \in \mathcal{A}^c} \omega_k \geq 2c_3 n^{-\kappa}$.

Remark 1. [Condition 1\(a\)](#) is a minimum signal strength condition that assumes the squared projection correlations between active features and response are uniformly bounded below, and cannot converge to zero too fast as n diverges. [Condition 1\(b\)](#) imposes an assumption on the gap of signal strength between active and inactive features. [Condition 1\(a\)](#) is weaker than [Condition 1\(b\)](#), since ω_k is always nonnegative. Such minimum signal strength condition can be viewed as a sparsity assumption that guarantees active features to be distinguished from the inactive ones. [Condition 1](#) is a very mild condition as we allow the minimum signal strength to converge to zero as the sample size diverges. Empirically, one may verify [Condition 1](#) by plotting the sorted empirical projection correlations of features in descending order and visually identifying whether there exists an “elbow” shape, or whether the sequence can be segmented into two groups. A more rigorous verification approach is to consider a multiple-testing procedure. For each feature, one can calculate the projection correlation based test statistic proposed in [Zhu et al. \(2017\)](#). Since it is not the main focus of this article, we do not pursue more details.

The following two theorems state the sure screening property and the rank consistency property of PC-Screen.

Theorem 2 (Sure screening). Under [Condition 1\(a\)](#), take $\delta \leq \min_{k \in \mathcal{A}} \omega_k/2$, we have

$$\Pr(\mathcal{A} \subseteq \widehat{\mathcal{A}}(\delta)) \geq 1 - O(s \exp\{-c_4 n^{1-2\kappa}\}), \quad (6)$$

where c_4 is a positive constant.

In [Theorem 2](#), if we set $\delta = c_3 n^{-\kappa}$, which satisfies the condition $\delta \leq \min_{k \in \mathcal{A}} \omega_k/2$, we have

$$\Pr(\mathcal{A} \subseteq \widehat{\mathcal{A}}(c_3 n^{-\kappa})) \geq 1 - O(s \exp\{-c_4 n^{1-2\kappa}\}). \quad (7)$$

From (7), we know that if $\delta = c_3 n^{-\kappa}$, all active features are selected with probability approaching 1 as $n \rightarrow \infty$. In fact, any choice of $\delta \leq c_3 n^{-\kappa}$ leads to the sure screening property. With the same choice of δ , [Li, Zhong, and Zhu \(2012\)](#) showed that the distance correlation based screening method (DC-SIS) satisfies

$$\Pr(\mathcal{A} \subseteq \widehat{\mathcal{A}}(c_3 n^{-\kappa})) \geq 1 - O(s \exp\{-c'_4 n^{1-2(\kappa+\eta)}\} + n \exp\{-c''_4 n^\eta\}),$$

where c'_4, c''_4 , and η are positive constants. Thus, PC-Screen achieves a faster rate than the DC-SIS since (1) we do not have the term $n \exp\{-c''_4 n^\eta\}$ and (2) we do not have an extra η in the power of the first term. The faster rate of PC-Screen is due to the fact that projection correlation is not sensitive to dimensionality and free of moment conditions.

Theorem 3 (Rank consistency). Under [Condition 1\(b\)](#), we have

$$\Pr\left(\min_{k \in \mathcal{A}} \widehat{\omega}_k - \max_{k \in \mathcal{A}^c} \widehat{\omega}_k > 0\right) > 1 - O(p \exp\{-c_5 n^{1-2\kappa}\}),$$

where c_5 is some positive constant. If $\log p = o(n^{1-2\kappa})$ with $0 \leq \kappa < 1/2$, then we have

$$\liminf_{n \rightarrow \infty} \left(\min_{k \in \mathcal{A}} \widehat{\omega}_k - \max_{k \in \mathcal{A}^c} \widehat{\omega}_k \right) > 0, \text{ almost surely.}$$

The rank consistency in [Theorem 3](#) is a stronger result than the sure screening property. When the signal strength gap between active and inactive features satisfies [Condition 1\(b\)](#), the active features are always ranked ahead of the inactive ones with high probability. In other words, there exists a choice of δ on the solution path that can perfectly separate the active and inactive sets with high probability.

3. Screening With FDR Control

3.1. Motivation

In the PC-Screen procedure, the threshold δ controls the cut-off between active and inactive features. [Theorem 2](#) suggests choosing $\delta = cn^{-\kappa}$ for some positive constants c and $\kappa < 1/2$. With certain model assumptions, the threshold δ (or equivalently c and κ) can be selected by cross-validation or information criterion approaches. However, in the model-free setup, such approaches are not directly applicable as the loss functions that measure the goodness of fit are not well defined. More recently, [Zhao and Li \(2012\)](#) and [Song et al. \(2014\)](#) suggested to select the threshold by controlling the false positive rate and estimating the proportion of active features s/p , respectively. These two

methods are tailored for survival analysis and are not directly applicable to model-free setup.

In practice, one can arbitrarily select a conservative threshold to ensure that all active features are included with high probability. However, it will include too many inactive features and inflate the FDR. Therefore, selecting the threshold parameter for model-free screening methods will inevitably cause the issue of balancing the trade-off between the sure screening property and FDR control. In this section, we propose a two-step procedure to address this issue by using knockoff features. The construction of knockoff features does not depend on the regression model, and hence is suitable for the model-free setting. Also, the second-order knockoff features can be easily computed by an equicorrelated construction or solving a semidefinite program (Barber and Candès 2015). The proposed procedure enjoys sure screening property and controls FDR simultaneously with high probability.

3.2. Knockoff Features: A Brief Review

Recently, knockoff has drawn huge attention due to its success in variable selection and many important applications such as genome-wide association studies. The concept of knockoff was first proposed in Barber and Candès (2015) for the fixed design matrix problem and then extended to the random design matrix setting known as Model-X knockoffs (Candès et al. 2018). For a more detailed development of knockoff, see Barber and Candès (2015), Candès et al. (2018), Fan, Demirkaya, Li and Lv (2020), and references therein. In this subsection, we briefly review the notations and definitions of knockoffs for further discussions.

Let $\mathbf{y} = (Y_1, \dots, Y_q)^T$ be a response vector of q variables and $\mathbf{x} = (X_1, \dots, X_p)^T$ be a covariate of p features. We say $\tilde{\mathbf{x}} = (\tilde{X}_1, \dots, \tilde{X}_p)^T$ is a knockoff copy of \mathbf{x} if it satisfies the following conditions.

Condition 2 (Exact knockoff features).

- (a) Swap X_j with its knockoff counterpart \tilde{X}_j does not change the joint distribution of $(\mathbf{x}, \tilde{\mathbf{x}})$ for $j = 1, \dots, p$;
- (b) Given $\mathbf{x}, \tilde{\mathbf{x}}$ is independent of \mathbf{y} , that is, $\tilde{\mathbf{x}} \perp\!\!\!\perp \mathbf{y} | \mathbf{x}$.

Remark 2. Condition 2(a) requires the original and knockoff features to be pairwise exchangeable. Condition 2(b) indicates that knockoff features are conditionally independent of response variables; this is trivially satisfied if $\tilde{\mathbf{x}}$ is generated without using the information of \mathbf{y} .

Constructing knockoff features that exactly follow Condition 2 is challenging, especially when the dimensionality of features p is large. In general, generating exact knockoff features requires the knowledge of the underlying distribution of \mathbf{x} , which is usually not available in practice. Barber and Candès (2015) studied the variable selection problem with knockoffs for fixed design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$. The construction of exact knockoffs is not necessary if assuming the response \mathbf{y} follows a linear regression model with Gaussian error and $n \geq 2p$. A more recent study (Candès et al. 2018) proposed to construct exact knockoff features in a manner of generating sequential

conditional independent pairs under the assumption that the distribution of \mathbf{x} is known.

Without the knowledge of the distribution of \mathbf{x} , one can construct approximate second-order knockoff features such that $(\mathbf{x}, \tilde{\mathbf{x}})$ is pairwise exchangeable with respect to the first two moments. In other words, the mean vector and covariance matrix of $(\mathbf{x}, \tilde{\mathbf{x}})$ is invariant if we swap X_j and \tilde{X}_j for any $j = 1, \dots, p$. The invariant of mean is trivial, and can be achieved by forcing $E(\mathbf{x}) = E(\tilde{\mathbf{x}})$. Suppose $\text{cov}(\mathbf{x}) = \Sigma$, the second-order pairwise exchangeable condition is equivalent to

$$\text{cov}(\mathbf{x}, \tilde{\mathbf{x}}) = \mathbf{G}, \quad \text{where } \mathbf{G} = \begin{bmatrix} \Sigma & \Sigma - \text{diag}\{\mathbf{h}\} \\ \Sigma - \text{diag}\{\mathbf{h}\} & \Sigma \end{bmatrix}, \quad (8)$$

where $\mathbf{h} = (h_1, \dots, h_p)^T$ is a vector that makes \mathbf{G} a positive semidefinite covariance matrix.

Barber and Candès (2015) introduced two approaches to construct the second-order knockoffs. The first approach is known as the *equicorrelated construction*, which sets

$$h_j = 2\lambda_{\min}(\Sigma) \vee 1 \quad \text{for } j = 1, \dots, p, \quad (9)$$

where $\lambda_{\min}(\Sigma)$ is the minimum eigenvalue of Σ , and $a \vee b$ denotes the larger one between a and b . The second method, named *semidefinite program*, finds h_j by solving a semidefinite program of the following form

$$\begin{aligned} & \text{minimize} \quad \sum_j |1 - h_j|, \\ & \text{subject to} \quad h_j \geq 0, \text{diag}\{\mathbf{h}\} \preceq 2\Sigma. \end{aligned} \quad (10)$$

However, both methods are not directly applicable in the high-dimensional scenarios since both (9) and (10) require $2p < n$.

Remark 3. When $(\mathbf{x}, \tilde{\mathbf{x}})$ is jointly Gaussian, the equivalence of the first two moments implies the equivalence of the joint distribution and hence (8) constructs exact knockoff features. However, when the Gaussian assumption does not hold, the accuracy of the second-order approximation depends on the impact of ignoring higher order moments of \mathbf{x} . Also, the covariance matrix Σ is usually unknown and needs to be estimated. Hence, the estimation accuracy of Σ also affects the validation of the second-order approximation of exact knockoff features. In addition to sample covariance estimator, more sophisticated estimators of Σ can be obtained under additional structure or moment conditions, see Fan, Liao, and Mincheva (2013) and Ke et al. (2019), among others. As the above two issues are not of key interest in this article, we do not pursue further in these directions.

3.3. FDR Control With Knockoff Features

Suppose $\tilde{\mathbf{x}} = (\tilde{X}_1, \dots, \tilde{X}_p)^T$ is a knockoff copy of $\mathbf{x} = (X_1, \dots, X_p)^T$, we propose to measure the population level dependence between X_j and response \mathbf{y} by the following quantity

$$W_j = \text{PC}(X_j, \mathbf{y})^2 - \text{PC}(\tilde{X}_j, \mathbf{y})^2, \quad j = 1, \dots, p, \quad (11)$$

where $\text{PC}(X_j, \mathbf{y})$ and $\text{PC}(\tilde{X}_j, \mathbf{y})$ are the projection correlations as defined in (2). When \tilde{X}_j is an exact knockoff feature of X_j ,

W_j will be a nonnegative quantity. Further, $W_j > 0$ implies the distribution of \mathbf{y} depends on X_j and $W_j = 0$ if \mathbf{y} is independent with X_j conditioning on active features.

Given a random sample $\{\mathbf{Y}, \mathbf{X}, \tilde{\mathbf{X}}\}$ drawn from $\{\mathbf{y}, \mathbf{x}, \tilde{\mathbf{x}}\}$, we estimate W_j by

$$\hat{W}_j = \widehat{\text{PC}}(\mathbf{X}_j, \mathbf{Y})^2 - \widehat{\text{PC}}(\tilde{\mathbf{X}}_j, \mathbf{Y})^2, \quad j = 1, \dots, p. \quad (12)$$

where $\widehat{\text{PC}}(\mathbf{X}_j, \mathbf{Y})$ and $\widehat{\text{PC}}(\tilde{\mathbf{X}}_j, \mathbf{Y})$ are sample projection correlations defined in (5). Intuitively, a large positive value of \hat{W}_j provides some evidence that the distribution of \mathbf{y} depends on X_j . On the other hand, if X_j is an inactive feature, $|\hat{W}_j|$ is likely be small and \hat{W}_j is equally likely to be positive or negative. This intuition is justified by the following lemma.

Lemma 1. Let $\tilde{\mathbf{x}}$ be an exact knockoff copy of \mathbf{x} and $\mathcal{A}^c = \{j_1, \dots, j_r\}$. Then

- (i) $W_{j_k} = 0$ for all $j_k \in \mathcal{A}^c$.
- (ii) Conditioning on $|\hat{\mathbf{w}}| = (|\hat{W}_1|, \dots, |\hat{W}_p|)^T$, I_{j_1}, \dots, I_{j_r} follow iid Bernoulli(0.5), where $I_{j_k} = 1$ if $\hat{W}_{j_k} > 0$ and 0 otherwise.

For a fixed threshold $t > 0$, the false discovery proportion (FDP) is defined as

$$\text{FDP}(t) = \frac{\#\{j \in \mathcal{A}^c : \hat{W}_j \geq t\}}{\#\{j : \hat{W}_j \geq t\}},$$

where $\#\{\cdot\}$ is the cardinality of a set and we follow the convention that $0/0 = 0$. The FDR is defined as the expectation of FDP, that is, $\text{FDR}(t) = \text{E}[\text{FDP}(t)]$. According to Lemma 1, \hat{W}_j is equally likely to be positive or negative if X_j is an inactive feature. Therefore, we have

$$\#\{j \in \mathcal{A}^c : \hat{W}_j \geq t\} \approx \#\{j \in \mathcal{A}^c : \hat{W}_j \leq -t\} \leq \#\{j : \hat{W}_j \leq -t\},$$

which leads to a conservative estimation of $\text{FDP}(t)$,

$$\widehat{\text{FDP}}(t) = \frac{\#\{j : \hat{W}_j \leq -t\}}{\#\{j : \hat{W}_j \geq t\}}.$$

To control FDR at a prespecified level α , we follow the knockoff+ procedure (Barber and Candès 2015) to choose the threshold T_α as

$$T_\alpha = \min \left\{ t \in \mathcal{W} : \frac{1 + \#\{j : \hat{W}_j \leq -t\}}{\#\{j : \hat{W}_j \geq t\}} \leq \alpha \right\}, \quad (13)$$

where $\mathcal{W} = \{|\hat{W}_j| : 1 \leq j \leq p\}/\{0\}$. The extra term 1 in the numerator makes the choice of T_α slightly more conservative. Then, the active set is selected as

$$\widehat{\mathcal{A}}(T_\alpha) = \{j : \hat{W}_j \geq T_\alpha, 1 \leq j \leq p\}. \quad (14)$$

3.4. PC-Knockoff Procedure

The knockoff feature construction methods discussed in Section 3.2 require $2p < n$ and hence are not applicable to high-dimensional scenarios. To address this issue, we propose a two-step procedure, named PC-Knockoff, to screen active features and control the FDR. To avoid mathematical challenges caused by the reuse of sample, we follow the simple sample splitting idea, which has been widely used in statistics and recent examples include hypothesis testing (e.g., Fan et al. 2019), error variance estimation (e.g., Chen, Fan, and Li 2018), variable selection (e.g., Barber and Candès 2019), large scale inference (e.g., Fan, Demirkaya, Li and Lv 2020), and so on. We partition the full sample into two disjoint subsamples with sample sizes n_1 and $n_2 = n - n_1$. More specifically, let $\mathbf{X}^{(1)} \in \mathbb{R}^{n_1 \times p}$ and $\mathbf{X}^{(2)} \in \mathbb{R}^{n_2 \times p}$ be a random partition of \mathbf{X} , and let \mathbf{Y} follow the same partition. Without loss of generality, we write

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}^{(1)} \\ \mathbf{X}^{(2)} \end{bmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{bmatrix} \mathbf{Y}^{(1)} \\ \mathbf{Y}^{(2)} \end{bmatrix}.$$

The two steps of PC-Knockoff procedure are introduced as follows:

1. Screening step: We rank all p features in descending order based on the sample projection correlation $\widehat{\text{PC}}(\mathbf{X}_j^{(1)}, \mathbf{Y}^{(1)})$. Then, we select the top d features such that $2d < n_2$. Denote the set of selected d features by $\widehat{\mathcal{A}}_1$. In practice, one can set d to be a relatively large value as long as it satisfies $2d < n_2$.

2. Knockoff step: Let

$$\mathbf{X}^{(2)} = \left(\mathbf{X}_{\widehat{\mathcal{A}}_1}^{(2)}, \mathbf{X}_{\widehat{\mathcal{A}}_1^c}^{(2)} \right).$$

We construct knockoff features for $\mathbf{X}_{\widehat{\mathcal{A}}_1}^{(2)}$ by either the *equicorrelated construction* as in (9) or the *semidefinite program* as in (10). Denoted $\tilde{\mathbf{X}}_{\widehat{\mathcal{A}}_1}^{(2)}$ the constructed knockoff features for $\mathbf{X}_{\widehat{\mathcal{A}}_1}^{(2)}$. Then, we calculate

$$\hat{W}_j = \widehat{\text{PC}}(\mathbf{X}_{\widehat{\mathcal{A}}_1, j}^{(2)}, \mathbf{Y}^{(2)})^2 - \widehat{\text{PC}}(\tilde{\mathbf{X}}_{\widehat{\mathcal{A}}_1, j}^{(2)}, \mathbf{Y}^{(2)})^2, \quad j = 1, \dots, d, \quad (15)$$

where $\mathbf{X}_{\widehat{\mathcal{A}}_1, j}^{(2)}$ and $\tilde{\mathbf{X}}_{\widehat{\mathcal{A}}_1, j}^{(2)}$ are the j th columns of $\mathbf{X}_{\widehat{\mathcal{A}}_1}^{(2)}$ and $\tilde{\mathbf{X}}_{\widehat{\mathcal{A}}_1}^{(2)}$, respectively. For a prespecified FDR level α , we use (13) to choose the threshold T_α and the set of selected active features is given by

$$\widehat{\mathcal{A}}(T_\alpha) = \{j : j \in \widehat{\mathcal{A}}_1, \hat{W}_j \geq T_\alpha\}.$$

The proposed two-step approach has two advantages. First, we apply the PC-Screen procedure to reduce the number of features from p to d , which allows us to construct second-order knockoff features. Second, by ruling out $p - d$ inactive features in the screening step, we reduce the total computation cost from $O(n^3 p)$ to $O(n_1^3 p + n_2^3 d)$. We summarize the PC-Knockoff procedure in Algorithm 1. In fact, the Algorithm 1 provides a general framework for feature screening with FDR control. One can easily modify Algorithm 1 by replacing PC with other measurement statistics such as Pearson correlation, distance correlation, etc.

Algorithm 1 PC-Knockoff

Input: $(\mathbf{X}, \mathbf{Y}) \in \mathbb{R}^{n \times p} \times \mathbb{R}^{n \times q}$, α , n_1 , and $d < n_2/2$, where $n_2 = n - n_1$.

Partition (\mathbf{X}, \mathbf{Y}) into $(\mathbf{X}^{(1)}, \mathbf{Y}^{(1)}) \in \mathbb{R}^{n_1 \times p} \times \mathbb{R}^{n_1 \times q}$ and $(\mathbf{X}^{(2)}, \mathbf{Y}^{(2)}) \in \mathbb{R}^{n_2 \times p} \times \mathbb{R}^{n_2 \times q}$.

1. Screening step

For $j = 1, \dots, p$, compute the sample squared PC: $\hat{\omega}_j^{(1)} = \widehat{\text{PC}}(\mathbf{X}_j^{(1)}, \mathbf{Y}^{(1)})^2$.

Select the top d features, that is, $\hat{\mathcal{A}}_1 = \{j : \hat{\omega}_j^{(1)} \text{ is among the largest } d\}$.

2. Knockoff step

Construct second-order knockoff features $\tilde{\mathbf{X}}_{\hat{\mathcal{A}}_1}^{(2)}$ for $\mathbf{X}_{\hat{\mathcal{A}}_1}^{(2)}$ using (9) or (10).

For all $j \in \hat{\mathcal{A}}_1$, compute $\hat{W}_j = \widehat{\text{PC}}(\mathbf{X}_{\hat{\mathcal{A}}_1, j}^{(2)}, \mathbf{Y}^{(2)})^2 - \widehat{\text{PC}}(\tilde{\mathbf{X}}_{\hat{\mathcal{A}}_1, j}^{(2)}, \mathbf{Y}^{(2)})^2$.

Choose the threshold T_α by solving (13).

Output: Set of selected active features $\hat{\mathcal{A}}(T_\alpha) = \{j : j \in \hat{\mathcal{A}}_1, \hat{W}_j \geq T_\alpha\}$.

Remark 4. We want to add a note that **Algorithm 1** is not tuning free as one needs to specify the subsample size n_1 and target dimension d for the screening step. However, **Algorithm 1** is not sensitive to the choice of these two hyper-parameters as **Theorem 2** guarantees the sure screening property of the screening step under mild conditions. In practice, we suggest small n_1 . As a result, we leave a relatively large subsample for the knockoff step. A relatively large n_2 allows more features to be selected in the screening step (larger d) and more accurate second-order knockoff features constructed in the knockoff step.

Denote \mathcal{E} the event that the sure screening property is satisfied in the screening step, that is,

$$\mathcal{E} = \{\text{All active features are selected in the screening step}\}.$$

The sure screening property in **Theorem 2** ensures that, with a relatively large choice of d , the event \mathcal{E} holds with high probability. To be specific, let $\hat{\omega}_{(1)}^{(1)} \geq \hat{\omega}_{(2)}^{(1)} \geq \dots \geq \hat{\omega}_{(p)}^{(1)}$ be the order statistics of sample squared projection correlations based on $(\mathbf{X}^{(1)}, \mathbf{Y}^{(1)})$. If **Condition 1(a)** holds and $\hat{\omega}_{(d)}^{(1)} \leq c_3 n_1^{-\kappa}$, then event \mathcal{E} holds with probability at least $1 - O(s \exp\{-c_4 n_1^{1-2\kappa}\})$. Conditioning on \mathcal{E} , the following theorem states that the PC-Knockoff procedure can control the FDR of selected features under the prespecified level of α .

Theorem 4. Let $\tilde{\mathbf{X}}_{\hat{\mathcal{A}}_1}$ be a knockoff copy of $\mathbf{X}_{\hat{\mathcal{A}}_1}$ satisfying **Condition 2**. For any $\alpha \in [0, 1]$, the set of selected features $\hat{\mathcal{A}}(T_\alpha)$ given by **Algorithm 1** satisfies

$$\text{FDR} = \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{A}^c \cap \hat{\mathcal{A}}(T_\alpha)\}}{\#\{j : j \in \hat{\mathcal{A}}(T_\alpha)\} \vee 1} \mid \mathcal{E} \right] \leq \alpha.$$

Theorem 4 states that, with exact knockoff features and conditioning on \mathcal{E} , the PC-Knockoff procedure can control the FDR under the prespecified level $\alpha \in [0, 1]$. As \mathcal{E} occurs with

probability close to 1, the FDR can be controlled even without conditioning on \mathcal{E} (Barber and Candès 2019). We refer to Barber and Candès (2019) and Fan, Demirkaya, Li and Lv (2020) for more discussion regarding FDR control in two-step procedures.

We hope the PC-Knockoff procedure can maintain sure screening property and control FDR simultaneously. This task is challenging as the procedure needs to balance the trade-off between Type I and Type II errors. To guarantee the sure screening property, the procedure is likely to select an over-fitted model which leads to higher Type I errors as well as FDR. On the other hand, the FDR control forces a parsimonious but possibly under-fitted model that can increase Type II errors and ruin the sure screening property. In the following, we study the conditions under which this challenging task can be achieved by the PC-Knockoff procedure. Conditioning on \mathcal{E} , the following theorem states that the simultaneous achievement of sure screening property and FDR control under level α , which depends on the relationship between s and α .

Theorem 5. Under the conditions of **Theorem 4** and further assume $\min_{k \in \mathcal{A}} W_j \geq 4c_3 n_2^{-\kappa}$ for some $c_3 > 0$ and $0 \leq \kappa < 1/2$.

- (i) If $\alpha \geq 1/s$, we have $\Pr(\mathcal{A} \subseteq \hat{\mathcal{A}}(T_\alpha) \mid \mathcal{E}) \geq 1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$.
- (ii) If $\alpha < 1/s$, we have $\Pr(\mathcal{A} \subseteq \hat{\mathcal{A}}(T_\alpha) \cup \hat{\mathcal{A}}(T_\alpha) = \emptyset \mid \mathcal{E}) \geq 1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$. Furthermore, if $s > 2$, we have $\Pr(\mathcal{A} \subseteq \hat{\mathcal{A}}(T_\alpha) \mid \mathcal{E}) \leq C(s) + O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$, where $0 < C(s) < 1$ is a constant that only depends on s .

The part (i) of **Theorem 5** together with **Theorem 4** suggest that if α is chosen to be greater or equal to $1/s$, the PC-Knockoff procedure enjoys the sure screening property and controls FDR under α with high probability. When α is chosen to be smaller than $1/s$, we either recover the active set or end up with an empty set with high probability. The probability of recovering the active set is upper bounded by some constant $C(s)$ depending on s . Therefore, there is no guarantee that PC-Knockoff can select all active features while controlling FDR under α . As we know, the value of α controls the amount of Type I errors that we can tolerate. One can imagine that the smaller the α is, the more challenging the task is to achieve sure screening and FDR control simultaneously. It is because we allow fewer and fewer inactive features to be selected. There is a phase transition that happens when α goes below $1/s$. To satisfy the sure screening property, the procedure will fail to control the FDR at α with nonnegligible probability if any inactive feature is selected. In other words, the procedure has to exactly recover the true active set to satisfy FDR control and sure screening simultaneously. We numerically validate this phase transition phenomenon through a simulated example in the supplementary materials.

Theorem 5 discourages us to pursue a too aggressive α in practice as the PC-Knockoff procedure may lose the sure screening property. The phase transition between part (i) and part (ii) can also be used as a rule of thumb guideline to estimate s in practice. For a sequence of grid points of α in $(0, 1)$, we find the largest grid point α^* such that the PC-Knockoff procedure selects an empty set. Then, we can roughly estimate \hat{s} as the integer part of $1/\alpha^*$.

Remark 5. As a two-step procedure, the power of the PC-Knockoff is conditioned on the event that the sure screening property is satisfied in the screening step, that is, \mathcal{E} . This together with the results in [Theorem 5](#) yield that the probability that the PC-Knockoff procedure does not lose any power is at least $1 - O(s \exp\{-c_4 n_1^{1-2\kappa}\}) - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$. Empirically, one may follow the “data recycling” idea proposed in [Barber and Candès \(2019\)](#) to improve the power of the PC-Knockoff procedure.

As a byproduct of the PC-Knockoff procedure, the statistic \widehat{W}_j defined in [\(12\)](#) can also be used to measure the marginal dependence between the response and the j th feature. Given a threshold δ , we can screen a model based on \widehat{W}_j by considering the following set

$$\widehat{\mathcal{A}}_W(\delta) = \{j : \widehat{W}_j \geq \delta, 1 \leq j \leq p\}.$$

The next theorem states that the screening procedure based on \widehat{W}_j also enjoys the sure screening and rank consistency properties.

Theorem 6. Suppose $\widetilde{\mathbf{X}}$ is an exact knockoff copy for \mathbf{X} and $\min_{k \in \mathcal{A}} W_j \geq 2c_3 n^{-\kappa}$ for some positive constants $c_3 > 0$ and $0 < \kappa < 1/2$, then

- (i) $\Pr(\mathcal{A} \subseteq \widehat{\mathcal{A}}_W(\delta)) \geq 1 - O(s \exp\{-c_4 n^{1-2\kappa}\})$ for $\delta \leq c_3 n^{-\kappa}$.
- (ii) $\Pr\left(\min_{k \in \mathcal{A}} \widehat{W}_k - \max_{k \in \mathcal{A}^c} \widehat{W}_k > 0\right) \geq 1 - O(p \exp\{-c_4 n^{1-2\kappa}\})$.

Recall that the rank consistency property in [Theorem 3](#) requires a minimum signal gap between active and inactive sets, that is, $\min_{k \in \mathcal{A}} \omega_k - \max_{k \in \mathcal{A}^c} \omega_k > 2c_3 n^{-\kappa}$. However, the rank consistency result in [Theorem 6](#) only requires a minimum signal strength of active features and no condition is imposed on the inactive set. Due to the construction of W_j , signals of inactive features are canceled out by their knockoff counterparts. As a result, the rank consistency property holds even when some inactive features are spuriously correlated with the response. If we can construct high-quality knockoff features, the screening procedure based on \widehat{W}_j can be more powerful than PC-Screen.

4. Numerical Examples

4.1. Screening Performance

In this subsection, we use simulated examples to assess the finite sample performance of the proposed projection correlation based feature screening procedure (PC-Screen) and compare it with SIS ([Fan and Lv 2008](#)), DC-SIS ([Li, Zhong, and Zhu 2012](#)), and bias-corrected distance correlation based screening (bcDC-SIS, [Székely and Rizzo 2014](#)). Within each replication, we rank the features in descending order by the above four screening criteria and record the minimum model size that contains all active features. The screening performance is measured by the 5%, 25%, 50%, 75%, and 95% quantiles of the minimum model size over 200 replications. Throughout this subsection, we denote $\Sigma = (\sigma_{ij})_{p \times p}$ with $\sigma_{ij} = 0.5^{|i-j|}$. To mimic ultrahigh-dimensional scenario, we set $n = 100$ and $p = 5000, 10000$ for each example.

4.1.1. Example 1: Linear and Poisson Models

Consider the linear model $Y = \mathbf{x}^T \boldsymbol{\beta} + \varepsilon$ with $\boldsymbol{\beta} = (1_5^T, \mathbf{0}_{p-5}^T)^T$. We generate covariates \mathbf{x} and ε independently from the following four scenarios.

Model 1.a: $\mathbf{x} \sim N(\mathbf{0}, \Sigma)$ and $\varepsilon \sim N(0, 1)$.

Model 1.b: $\mathbf{x} \sim N(\mathbf{0}, \Sigma)$ and $\varepsilon \sim \text{Cauchy}(0, 1)$.

Model 1.c: $\mathbf{u} \sim \text{Cauchy}(\mathbf{0}, \mathbf{I}_p)$, $\mathbf{x} = \Sigma^{1/2} \mathbf{u}$, and $\varepsilon \sim N(0, 1)$.

Model 1.d: $\mathbf{u} \sim \text{Cauchy}(\mathbf{0}, \mathbf{I}_p)$, $\mathbf{x} = \Sigma^{1/2} \mathbf{u}$, and $\varepsilon \sim \text{Cauchy}(0, 1)$.

In above models, $\text{Cauchy}(\mathbf{0}, \mathbf{I}_p)$ stands for the p -dimensional standard Cauchy distribution which is heavy-tailed. Hence, in Models 1.b–1.d, at least one of \mathbf{x} and ε is heavy-tailed. We also consider the following two Poisson regression models

Model 1.e: (Continuous) $Y = \exp\{\mathbf{x}^T \boldsymbol{\beta}\} + \varepsilon$, where $\varepsilon \sim N(0, 1)$.

Model 1.f: (Discrete) $Y \sim \text{Poisson}(\exp\{\mathbf{x}^T \boldsymbol{\beta}\})$.

Let $\boldsymbol{\beta} = (2_5^T, \mathbf{0}_{p-5}^T)^T$ and we draw \mathbf{x} from $N(\mathbf{0}, \Sigma)$. Model 1.e is the Poisson regression model with continuous response while the response in Model 1.f is discrete.

The quantiles of the minimum model size that includes all five active features are summarized in [Table 1](#). In the linear benchmark Model 1.a, all four competitors perform well. For Models 1.b–1.d, the SIS completely fails at the presence of heavy-tailed features and errors. Both distance correlation-based methods struggle to maintain a reasonable model size at 75% and 95% quantiles. In contrast, PC-Screen works reasonably well in all scenarios and outperforms the other three methods by big margins. Similarly, for the Poisson models 1.e and 1.f, PC-Screen can recover the true active set with a model size close to 5 while the other three methods can perform as bad as random guesses at 75% and 95% quantiles.

4.1.2. Example 2: Nonlinear Models

Consider the following four nonlinear data generating models

Model 2.a: $Y = 5X_1 + 2 \sin(\pi X_2/2) + 2X_3 \mathbf{1}\{X_3 > 0\} + 2 \exp\{5X_4\} + \varepsilon$.

Model 2.b: $Y = 3X_1 + 3X_2^3 + 3X_3^{-1} + 5\mathbf{1}\{X_4 > 0\} + \varepsilon$.

Model 2.c: $Y = 1 - 5(X_2 + X_3)^3 \exp\{-5(X_1 + X_4^2)\} + \varepsilon$.

Model 2.d: $Y = 1 - 5(X_2 + X_3)^{-3} \exp\{1 + 10 \sin(\pi X_1/2) + 5X_4\} + \varepsilon$.

Models 2.a and 2.b admit additive structure while Models 2.c and 2.d have more challenging nonlinear structures. In addition, we generate $\mathbf{x} \sim N(\mathbf{0}, \Sigma)$ and $\varepsilon \sim N(0, 1)$. Hence, for each model above, the active set contains the first four covariates in \mathbf{x} .

The quantiles of the minimum model size that includes all four active features are presented in [Table 2](#). Again, we observe that PC-Screen significantly outperforms the other three methods in all scenarios. For Model 2.a, the 50% quantile of the minimum model size of PC-Screen is exactly four while the other three methods need much larger model sizes to recover the active set. For Model 2.b, DC-SIS and bcDC-SIS preform comparably to PC-Screen at 5% and 25% quantiles but much worse for other higher quantiles. For Models 2.c and 2.d, PC-Screen performs reasonably well while the other methods fail to effectively screen out the inactive features. The 95% quantiles

Table 1. The quantiles of minimum model size for linear and Poisson models in Example 1 over 200 replications.

	Model 1.a					Model 1.b				
	5%	25%	50%	75%	95%	5%	25%	50%	75%	95%
<i>p</i> = 5000										
PC-Screen	5.0	5.0	5.0	5.0	8.0	5.0	5.0	6.0	14.0	125.0
DC-SIS	5.0	5.0	5.0	5.0	6.0	5.0	5.0	10.0	81.2	1305.0
bcDC-SIS	5.0	5.0	5.0	5.0	6.0	5.0	5.0	6.0	20.5	231.4
SIS	5.0	5.0	5.0	5.0	5.0	6.0	238.0	1833.0	3878.5	4915.0
<i>p</i> = 10,000										
PC-Screen	5.0	5.0	5.0	5.0	7.0	5.0	5.0	8.0	23.0	233.5
DC-SIS	5.0	5.0	5.0	5.0	6.0	5.0	6.8	21.0	204.2	3511.2
bcDC-SIS	5.0	5.0	5.0	5.0	6.0	5.0	5.0	10.0	43.2	718.8
SIS	5.0	5.0	5.0	5.0	5.0	16.0	703.2	3418.5	7432.0	9651.0
Model 1.c						Model 1.d				
	5%	25%	50%	75%	95%		5%	25%	50%	95%
<i>p</i> = 5000										
PC-Screen	5.0	5.0	6.0	8.0	50.9	5.0	5.0	6.0	10.2	139.9
DC-SIS	5.0	8.0	42.5	143.0	722.5	5.0	16.0	54.0	189.0	701.6
bcDC-SIS	5.0	5.0	6.0	14.2	80.4	5.0	5.0	8.0	21.8	156.8
SIS	5.0	39.0	81.5	374.0	3244.4	5.0	45.8	130.5	523.5	3241.0
<i>p</i> = 10,000										
PC-Screen	5.0	5.0	6.0	8.0	92.4	5.0	6.0	6.0	13.0	176.2
DC-SIS	5.0	18.2	78.0	277.5	1567.5	5.0	30.8	113.0	410.0	2036.5
bcDC-SIS	5.0	5.0	7.0	19.2	179.1	5.0	5.0	10.0	27.0	412.6
SIS	6.0	58.8	180.5	773.2	4189.0	8.9	73.0	244.5	959.8	5777.7
Model 1.e						Model 1.f				
	5%	25%	50%	75%	95%		5%	25%	50%	95%
<i>p</i> = 5000										
PC-Screen	5.0	5.0	5.0	5.0	17.2	5.0	5.0	5.0	5.0	7.0
DC-SIS	90.3	396.0	897.5	1762.5	3787.3	76.3	396.0	893.5	1764.8	3471.8
bcDC-SIS	22.7	82.2	259.5	669.2	2358.7	15.0	87.2	266.5	821.2	2471.5
SIS	178.6	604.8	1137.0	2319.8	4303.4	186.0	606.2	1210.5	2253.0	4261.6
<i>p</i> = 10,000										
PC-Screen	5.0	5.0	5.0	6.0	23.0	5.0	5.0	5.0	5.0	12.0
DC-SIS	138.0	788.8	1878.5	3301.0	6831.9	154.2	729.2	1725.5	3424.0	6832.8
bcDC-SIS	45.7	163.5	534.5	1520.5	5415.2	30.0	175.8	509.0	1513.2	5580.0
SIS	462.8	1276.8	2460.0	4164.5	8622.5	512.6	1271.2	2484.5	4281.0	8416.3

NOTE: The true model size is 5.

of SIS, DC-SIS, and bcDC-SIS are almost as large as *p*. This indicates, in the worst-case scenario, SIS, DC-SIS, and bcDC-SIS are hopeless to effectively reduce the dimensionality without missing any active feature.

4.1.3. Example 3: Multivariate Response Models

In this experiment, we investigate the performance of PC-Screen for multivariate response models. We omit SIS method in this experiment as it is not applicable to multivariate response problems. We generate $\mathbf{y} = (Y_1, Y_2)^T$ from a bivariate normal distribution with conditional mean $\mu_{Y|X} = (\mu_1(\mathbf{x}), \mu_2(\mathbf{x}))^T$ and covariance matrix $\Sigma_{Y|X} = (\sigma_{ij})_{2 \times 2}$, where $\sigma_{11} = \sigma_{22} = 1$ and $\sigma_{12} = \sigma_{21} = \sigma(\mathbf{x})$. Following the setting in Li, Zhong, and Zhu (2012), we set $\boldsymbol{\theta} = (2_4^T, \mathbf{0}_{p-4}^T)^T$ and generate $\mu_1(\mathbf{x})$, $\mu_2(\mathbf{x})$, and $\sigma(\mathbf{x})$ from the two models below.

Model 3.a: $\mu_1(\mathbf{x}) = \exp\{2(X_1 + X_2)\}$, $\mu_2(\mathbf{x}) = X_3 + X_4$, and $\sigma(\mathbf{x}) = \sin(\mathbf{x}^T \boldsymbol{\theta})$.

Model 3.b: $\mu_1(\mathbf{x}) = 2 \sin(\pi X_1/2) + X_3 + \exp\{1 + X_4\}$, $\mu_2(\mathbf{x}) = X_1^{-2} + X_2$, and $\sigma(\mathbf{x}) = (\exp\{\mathbf{x}^T \boldsymbol{\theta}\} - 1)/(\exp\{\mathbf{x}^T \boldsymbol{\theta}\} + 1)$.

The union of the active sets of $\mu_1(\mathbf{x})$, $\mu_2(\mathbf{x})$, and $\sigma_1(\mathbf{x})$ contains the first four covariates in \mathbf{x} . The simulation results are summarized in Table 3. Again, the PC-Screen method performs strikingly well compared to the other two methods.

4.2. FDR Control Performance

In this subsection, we use simulated examples to numerically assess the FDR control as well as sure screening property of the proposed PC-Knockoff procedure. We refer to Algorithm 1 for implementation details.

4.2.1. Example 4: FDR Control for Linear and Poisson Models

Consider the following five regression models with ten active variables.

Model 4.a: Same as Model 1.a except that $\boldsymbol{\beta} = (\mathbf{1}_s, \mathbf{0}_{p-s})$ with $s = 10$.

Model 4.b: Same as Model 4.a except that $\varepsilon \sim t_2$, the *t* distribution with degrees of freedom 2.

Table 2. The quantiles of minimum model size for nonlinear models in Example 2 over 200 replications.

	Model 2.a					Model 2.b				
	5%	25%	50%	75%	95%	5%	25%	50%	75%	95%
$p = 5000$										
PC-Screen	4.0	4.0	4.0	5.2	19.1	4.0	5.0	9.5	26.5	261.9
DC-SIS	600.0	1880.2	2994.5	4052.2	4701.3	4.0	7.8	61.0	541.5	2310.9
bcDC-SIS	488.4	1480.8	2863.0	3967.8	4855.9	4.0	6.0	21.5	88.0	893.8
SIS	709.0	2065.8	3062.5	4160.0	4869.4	54.0	658.5	2692.5	4213.0	4829.1
$p = 10,000$										
PC-Screen	4.0	4.0	4.0	5.0	31.0	4.0	5.8	13.0	48.8	393.9
DC-SIS	664.0	3162.5	5655.5	7605.8	9490.1	4.0	13.8	86.0	843.2	5529.2
bcDC-SIS	663.4	2605.8	5578.5	7745.2	9208.4	4.0	8.0	24.5	150.5	1403.9
SIS	986.2	4048.2	6068.0	8298.5	9752.6	64.5	1193.8	4488.0	8139.5	9725.6
Model 2.c										
	5%	25%	50%	75%	95%	Model 2.d				
	5%	25%	50%	75%	95%	5%	25%	50%	75%	95%
$p = 5000$										
PC-Screen	4.0	4.0	4.0	6.0	21.3	4.0	5.0	9.5	36.0	169.4
DC-SIS	397.6	1536.8	2750.5	3930.0	4721.8	1639.1	3138.5	3851.5	4434.8	4902.4
bcDC-SIS	196.5	1267.5	2774.0	4236.0	4986.2	605.5	1251.8	2073.5	2972.5	4209.1
SIS	421.2	1615.8	2920.0	4057.0	4761.0	2065.0	3487.8	4083.5	4659.0	4950.3
$p = 10,000$										
PC-Screen	4.0	4.0	4.0	7.2	25.0	4.0	5.0	13.0	62.5	297.4
DC-SIS	668.8	3628.2	6243.5	7920.5	9531.6	3409.6	6380.8	7705.0	8756.5	9790.6
bcDC-SIS	288.2	2006.5	4714.5	7370.2	9869.9	776.4	2567.8	4154.5	5517.8	8251.0
SIS	760.9	3609.5	6155.5	8129.8	9577.1	3333.4	6387.8	8007.5	9252.8	9847.2

NOTE: The true model size is 4.

Table 3. The quantiles of minimum model size for multivariate response models in Example 3 over 200 replications.

	Model 3.a					Model 3.b				
	5%	25%	50%	75%	95%	5%	25%	50%	75%	95%
$p = 5000$										
PC-Screen	4.0	4.0	4.0	4.0	6.0	4.0	4.0	6.0	18.0	114.4
DC-SIS	53.0	463.0	1211.0	2349.0	3774.2	641.1	2308.8	3307.0	4257.8	4838.0
bcDC-SIS	24.0	215.5	758.5	1965.0	3999.8	225.2	1270.2	2494.5	3596.8	4709.0
$p = 10,000$										
PC-Screen	4.0	4.0	4.0	4.0	9.0	4.0	4.0	8.0	30.2	302.9
DC-SIS	136.8	978.5	2237.5	4506.0	8106.6	1804.7	4510.5	6445.5	8153.5	9707.6
bcDC-SIS	83.0	546.0	1828.5	4264.5	7711.3	444.4	2217.0	4695.0	7122.5	9076.4

NOTE: The true model size is 4.

Model 4.c: Same as Model 4.a except that $\mathbf{x} = 0.9\mathbf{x}_1 + 0.1\mathbf{x}_2$, where \mathbf{x}_1 and \mathbf{x}_2 are independently drawn from $\mathbf{x}_1 \sim N(\mathbf{0}, \Sigma)$ and $\mathbf{x}_2 \sim t_2(\mathbf{0}, \Sigma)$, respectively.

Model 4.d: Same as Model 1.e except that $\boldsymbol{\beta} = 2 \cdot (\mathbf{1}_s, \mathbf{0}_{p-s})$ with $s = 10$.

Model 4.e: Same as Model 1.f except that $\boldsymbol{\beta} = 2 \cdot (\mathbf{1}_s, \mathbf{0}_{p-s})$ with $s = 10$.

In this example, we set $n = 1000$, $p = 5000$ and repeat 200 replications for each scenario. In each replication, we randomly divide the sample into two nonoverlapping subsamples. The sample size and target dimension in the screening step are set to be $n_1 = 250$ and $d = 100$, and the sample size used to construct knockoff features is $n_2 = 750$. In addition, the covariance matrix is set to be $\Sigma = (\sigma_{ij})$ with $\sigma_{ij} = 0.5^{|i-j|}$. The performance of FDR control is examined under a sequence of specified levels: $\alpha = 0.1, 0.15, 0.20, 0.25$, and 0.30 .

We summarize the results in Table 4 in which α is the prespecified FDR level, $|\widehat{\mathcal{A}}|$ is the average number of selected

variables, the column “ X_j ” represents the probability that the active variable X_j is selected, the column “All” represents the sure screening probability (i.e., the probability that all active variables are selected) and the column “ \widehat{FDR} ” is the empirical FDR (i.e., the average of empirical FDP).

According to Table 4, the proposed PC-Knockoff procedure controls the empirical FDR under the prespecified level α for almost all scenarios. The only exception is in Model 4.c we have $\widehat{FDR} = 0.254$ at $\alpha = 0.25$. In Model 4.c, \mathbf{x} is drawn from a mixture of multivariate normal and multivariate t_2 distributions. As a result, the second-order knockoffs may not approximate the exact ones very well. For the other four models, the second-order knockoffs perform well as features are normally distributed. Besides FDR control, the PC-Knockoff procedure maintains the sure screening property reasonably well in all scenarios. For Models 4.a–4.d, the probabilities of selecting all active variables are ranging from 91.5% to 99.5% regardless of the prespecified level α . Even for the more challenging scenario

Table 4. FDR control for linear and Poisson models in Example 4.

α	$ \hat{\mathcal{A}} $	X_1	X_2	X_3	X_4	X_5	X_6	X_7	X_8	X_9	X_{10}	All	$\widehat{\text{FDR}}$
Model 4.a													
0.10	11.245	0.995	0.995	1.000	0.995	0.995	0.995	0.995	0.995	0.990	0.990	0.997	
0.15	11.935	0.995	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.990	0.130	
0.20	12.860	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.188	
0.25	14.095	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.244	
0.30	15.270	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.285	
Model 4.b													
0.10	10.545	0.945	0.945	0.960	0.950	0.955	0.950	0.955	0.955	0.945	0.915	0.079	
0.15	11.350	0.985	0.990	0.995	0.995	0.995	1.000	0.995	1.000	0.990	0.970	0.920	0.100
0.20	12.460	0.990	0.995	1.000	0.995	0.995	1.000	0.995	1.000	0.985	0.955	0.164	
0.25	13.570	0.995	1.000	1.000	0.995	0.995	1.000	0.995	1.000	0.985	0.965	0.216	
0.30	14.705	1.000	1.000	1.000	0.995	0.995	1.000	1.000	1.000	0.990	0.980	0.260	
Model 4.c													
0.10	11.165	0.995	0.995	0.995	1.000	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.092
0.15	11.660	1.000	0.995	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.116	
0.20	13.075	1.000	0.995	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.193	
0.25	14.420	1.000	0.995	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.254	
0.30	15.570	1.000	0.995	1.000	1.000	1.000	1.000	1.000	1.000	0.995	0.995	0.299	
Model 4.d													
0.10	10.630	0.940	0.945	0.940	0.945	0.955	0.940	0.950	0.940	0.950	0.940	0.925	0.092
0.15	11.660	0.985	0.980	0.995	0.985	0.995	0.990	1.000	1.000	0.990	0.990	0.930	0.122
0.20	12.870	0.990	0.990	1.000	1.000	1.000	0.995	1.000	1.000	0.990	0.990	0.965	0.193
0.25	13.830	0.990	0.995	1.000	1.000	1.000	0.995	1.000	1.000	0.995	0.995	0.975	0.242
0.30	15.020	0.995	0.995	1.000	1.000	1.000	0.995	1.000	1.000	0.995	0.995	0.980	0.288
Model 4.e													
0.10	9.790	0.855	0.870	0.860	0.860	0.865	0.870	0.875	0.870	0.865	0.865	0.815	0.086
0.15	11.590	0.955	0.985	0.975	0.965	0.980	0.995	0.980	0.970	0.965	0.970	0.825	0.123
0.20	12.895	0.970	0.990	0.995	0.995	1.000	1.000	0.995	0.990	0.985	0.980	0.935	0.189
0.25	14.025	0.980	0.995	0.995	0.995	1.000	1.000	0.995	0.990	0.990	0.990	0.935	0.244
0.30	15.040	0.985	1.000	0.995	0.995	1.000	1.000	0.995	0.990	0.990	0.990	0.945	0.283

NOTE: The true model size is 10.

Table 5. Comparison between $\hat{\mathcal{A}}_1$ and $\hat{\mathcal{A}}(T_\alpha)$.

	PC-Screen			PC-Knockoff		
	$ \hat{\mathcal{A}}_1 $	All	$\widehat{\text{FDR}}$	$ \hat{\mathcal{A}}(T_\alpha) $	All	$\widehat{\text{FDR}}$
Model 4.a	100	0.995	0.90	12.9	0.995	0.188
Model 4.b	100	0.995	0.90	12.5	0.955	0.164
Model 4.c	100	0.995	0.90	13.0	0.995	0.193
Model 4.d	100	0.995	0.90	12.9	0.965	0.193
Model 4.e	100	0.995	0.90	12.9	0.935	0.189

NOTE: The number active features is $s = 10$ and prespecified FDR level for PC-Knockoff is $\alpha = 0.2$.

Model 4.e, the probability of selecting all active variables simultaneously is greater than 93.5% when $\alpha \geq 0.2$. Please notice that this is achieved under very small model sizes (i.e., small $|\hat{\mathcal{A}}|$), thanks to the knockoff features.

To illustrate how the two steps in PC-Knockoff work together, we also compare the active set $\hat{\mathcal{A}}_1$ selected by the screening step without knockoff features (i.e., selected by PC-Screen) and the active set $\hat{\mathcal{A}}(T_\alpha)$ selected by PC-Knockoff. Table 5 summarizes the average number of selected features $|\hat{\mathcal{A}}|$, the sure screening probability “All” and empirical FDR $\widehat{\text{FDR}}$ for Models 4.a–4.e. For the screening step, we conservatively select active features by including the top $d = 100$ features. The results show that $\hat{\mathcal{A}}_1$ enjoys the sure screening property but has a very high empirical FDR ($\widehat{\text{FDR}} = 0.90$). On the contrary, the active

set $\hat{\mathcal{A}}(T_\alpha)$ selected by the PC-Knockoff is able to control the FDR below the prespecified level with only a little sacrifice of power.

4.3. Supermarket Data

In this subsection, we apply the PC-Knockoff procedure to study a supermarket dataset (Wang 2009; Chen, Fan, and Li 2018). The dataset consists of $n = 464$ observations of daily records from a supermarket. The response Y is the number of customers visited the supermarket on that day. The covariates are sale volumes of $p = 6398$ products. Due to data privacy, the detailed product codes are not released in the dataset. Instead, we name the covariates by their indices, that is, X_1, \dots, X_{6398} . Both the response and predictors have been standardized to have zero mean and unit variance. The goal is to screen a parsimonious set of products whose sale volumes significantly contribute to the daily number of customers. Meanwhile, we want to control the FDR at level $\alpha = 0.1$.

The implementation of PC-Knockoff follows Algorithm 1. We set $n_1 = 200$ and $d = 50$ in the screening step. That is, we randomly choose 200 observations and use PC-Screen to prescreen 50 features in the screening step. Then we construct second-order knockoffs for the pre-screened 50 features using the remaining 264 observations in the knockoff step. Under the prespecified FDR level $\alpha =$

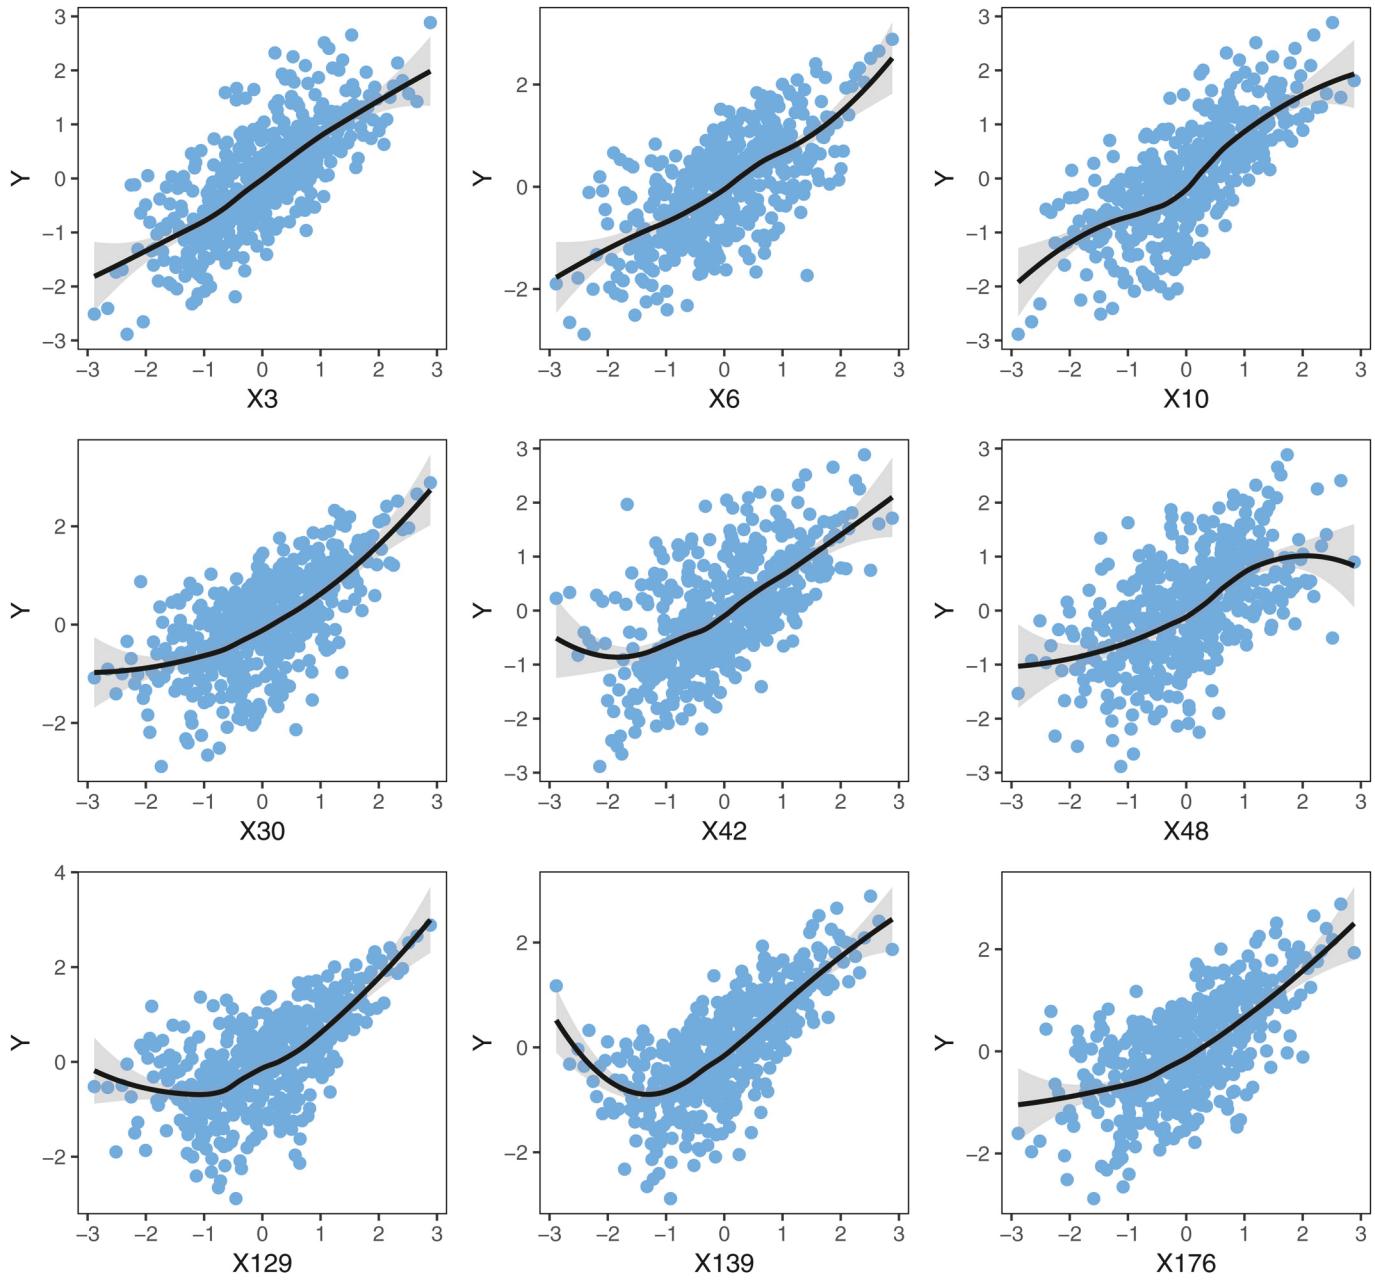


Figure 1. Scatterplots between the response and variables $X_3, X_6, X_{10}, X_{30}, X_{42}, X_{48}, X_{129}, X_{139}$, and X_{176} . The blue dots are observations. The black solid curves are fitted local polynomial regression curves. The gray shaded areas are corresponding confidence regions.

0.1, the PC-Knockoff procedure selects a model of 12 variables: $X_3, X_6, X_{10}, X_{11}, X_{30}, X_{42}, X_{48}, X_{71}, X_{129}, X_{139}, X_{176}$, and X_{400} . For each selected variable, we draw a scatterplot between this variable and the response. We observe obvious outliers in the scatterplots of X_{11}, X_{71} , and X_{400} . We report this observation in Figure 2, where red triangles denote potential outliers and blue dots denote the rest data points. Besides, red dashed curves and blue solid curves are the fitted local polynomial regression curves with and without potential outliers, respectively. The gray shaded areas are the 95% confidence regions. By comparing blue and red curves, we find the existence of outliers visually alters the fitted curves. This justifies the PC-Knockoff procedure is insensitive to the presence of outliers. The scatterplots between the response and other nine variables ($X_3, X_6, X_{10}, X_{30}, X_{42}, X_{48}, X_{129}, X_{139}$, and X_{176}) are presented in Figure 1, which includes

the local polynomial regression curves with the 95% confidence regions. Figure 1 indicates that PC-Knockoff can detect various functional relationships.

This supermarket dataset has also been analyzed by Chen, Fan, and Li (2018). They first apply DC-SIS to screen variables and then fit an additive model with selected variables. They further employ the Wald's χ^2 -test with the refitted cross-validation error variance estimate to determine if the selected features are significant at a prespecified level. As a result, Chen, Fan, and Li (2018) selected seven significant variables $X_3, X_6, X_{11}, X_{39}, X_{42}, X_{62}$, and X_{139} . Next, we compare the in-sample fitting and out-of-sample prediction performance between the model selected by PC-Knockoff and the one in Chen, Fan, and Li (2018) through a bootstrap experiment. In each replication, we randomly split the dataset into a training set

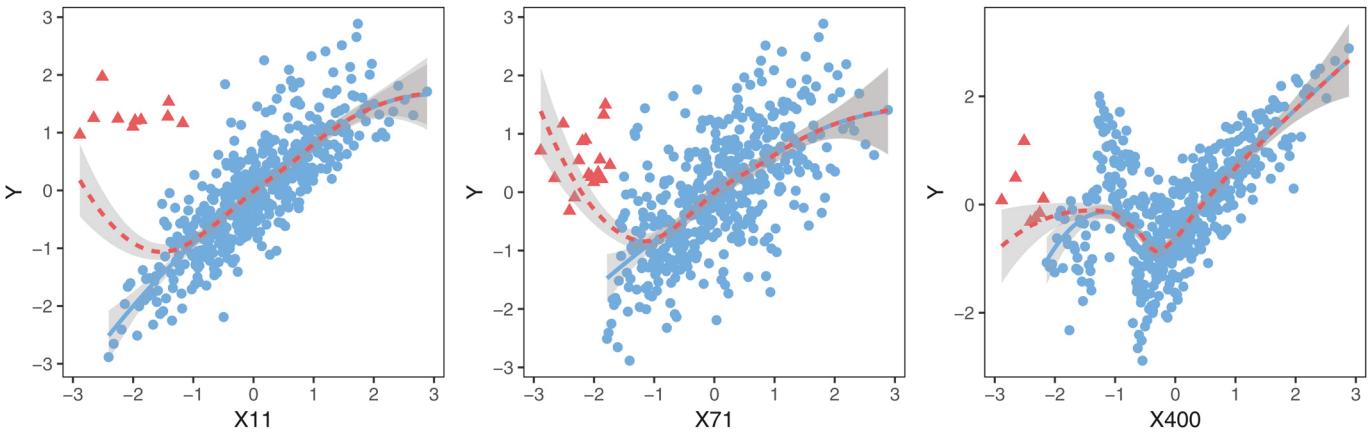


Figure 2. Scatterplots between the response and variables X_{11} , X_{71} , and X_{400} . The red triangles are the potential outliers and blue dots are the rest observations. The red dashed curves and blue solid curves are fitted local polynomial regression curves with and without the potential outliers, respectively. The gray shaded areas are corresponding confidence regions.

Table 6. Sample mean and sample standard deviation of R^2 for training set and test set over 200 bootstrap replications for supermarket data.

	Training set		Test set	
	Mean	SD	Mean	SD
PC-Knockoff	0.8681	0.0042	0.8608	0.0281
DC-SIS with χ^2 -test	0.8534	0.0043	0.8488	0.0276

of size 400 and a test set of size 64. We fit two additive models with the features selected by the two competitive methods, respectively. Then, we calculate and record the training and testing R^2 's for the two models. We repeat it for 200 replications. The sample mean and sample standard deviation of R^2 's are reported in Table 6, which show that the model selected by PC-Knockoff yields higher sample means of R^2 for both training set and test set.

Due to the limited space, we present an additional real application to microarray data in Section S.2 of the supplementary materials.

5. Conclusion

We study the feature screening problems under a general setup. The proposed PC-Screen method developed in Section 2 is novel in the sense that it does not impose any regression model assumption and is robust against possibly heavy-tailed data. The response variable is also allowed to be multidimensional. Regarding the theoretical aspect, the asymptotic results developed in Zhu et al. (2017) do not apply to the “large p small n problem,” and hence we develop nonasymptotic results for the empirical projection correlation. The theoretical analysis shows that the PC-Screen method satisfies sure screening and rank consistency properties. Numerically, we show PC-Screen outperforms popular competitors over various data generative processes. Moreover, this article tackles the FDR control problem in feature screening. In Section 3, we propose a two-step procedure named PC-Knockoff to determine the threshold of active features. With the sample splitting idea, we first screen the ultrahigh-dimensional features to a moderate model size and then further select a model with FDR control using a statistic

built upon knockoff features. The PC-Knockoff procedure controls FDR under a prespecified level while maintaining sure screening property with high probability. In practice, the PC-Knockoff procedure works well in various scenarios.

Appendix A: Proof of Results in Section 3

In the appendix, we provide the proofs of Theorems 4 and 5 in Section 3. We relegate the proofs of remaining results as well as some additional numerical examples to the supplementary materials.

A.1. Proof of Theorem 4

Denote $\widehat{\mathcal{A}}_1$, the set of features that are selected in the screening step. Throughout this proof, we restrict ourselves to the event $\mathcal{E} = \{\mathcal{A} \subset \widehat{\mathcal{A}}_1\}$. According to Theorem 2, the probability of \mathcal{E} is at least $1 - O(s \exp\{-c_4 n_1^{1-2K}\})$. To simplify the notations, we omit the condition on \mathcal{E} .

Denote by $\mathcal{B} = \mathcal{A} \cap \widehat{\mathcal{A}}_1$ and $\mathcal{B}^c = \mathcal{A}^c \cap \widehat{\mathcal{A}}_1$ the intersections between active and inactive sets with $\widehat{\mathcal{A}}_1$, respectively. Notice that any feature with $|\widehat{W}_j| = 0$ will not be included in the final selected set $\widehat{\mathcal{A}}(T_\alpha)$. Hence, without loss of generality, we assume $\widehat{\mathcal{A}}_1 = \{1, 2, \dots, d\}$ and $|\widehat{W}_1| \geq |\widehat{W}_2| \geq \dots \geq |\widehat{W}_d| > 0$. For ease of presentation, we further define $\widehat{W}_{d+1} = 0$. Then we have

$$\begin{aligned}
& \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \cap \widehat{\mathcal{A}}(T_\alpha)\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right] \\
&= \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \geq T_\alpha\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right] \\
&= \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \geq T_\alpha\}}{1 + \#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -T_\alpha\}} \right. \\
&\quad \left. \cdot \frac{1 + \#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -T_\alpha\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right] \\
&\leq \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \geq T_\alpha\}}{1 + \#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -T_\alpha\}} \right. \\
&\quad \left. \cdot \frac{1 + \#\{j : j \in \widehat{\mathcal{A}}_1 \text{ and } \widehat{W}_j \leq -T_\alpha\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right]
\end{aligned}$$

$$\leq \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \geq T_\alpha\}}{1 + \#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -T_\alpha\}} \cdot \alpha \right]. \quad (\text{A.1})$$

The first inequality holds since $\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -T_\alpha\} \subset \{j : j \in \widehat{\mathcal{A}}_1 \text{ and } \widehat{W}_j \leq -T_\alpha\}$ and the second inequality is due to the definition of T_α in (13). To find T_α , one can simply try different values of t starting from the smallest value $t = |\widehat{W}_{d+1}| = 0$, then move to the second smallest value $t = |\widehat{W}_d|$, then move to $t = |\widehat{W}_{d-1}|$, and so on. The procedure stops as soon as it finds a value of t satisfying (13). In this process, T_α can be regarded as a stopping time. More rigorously, for $k = d+1, d, d-1, \dots, 1$, we define

$$\begin{aligned} M(k) &= \frac{\#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \geq |\widehat{W}_k|\}}{1 + \#\{j : j \in \mathcal{B}^c \text{ and } \widehat{W}_j \leq -|\widehat{W}_k|\}} \\ &= \frac{\#\{j : j \in \mathcal{B}^c, j \leq k, \widehat{W}_j > 0\}}{1 + \#\{j : j \in \mathcal{B}^c, j \leq k, \widehat{W}_j \leq 0\}} \\ &=: \frac{V_+(k)}{1 + V_-(k)}, \end{aligned}$$

where $V_+(k) = \#\{j : j \in \mathcal{B}^c, j \leq k, \widehat{W}_j > 0\}$ and $V_-(k) = \#\{j : j \in \mathcal{B}^c, j \leq k, \widehat{W}_j \leq 0\}$. Let \mathcal{F}_k be the σ -algebra generated by $\{V_\pm(d+1), V_\pm(d), \dots, V_\pm(k), Z_{d+1}, Z_d, \dots, Z_k\}$ where $Z_{d+1} = 0$ and

$$Z_j = \begin{cases} 1 & \text{if } j \in \mathcal{B}, \\ 0 & \text{if } j \in \mathcal{B}^c. \end{cases}$$

As a result, given \mathcal{F}_k , we know whether k is in the active set \mathcal{B} or not.

Next we show that the process $M(d+1), M(d), \dots, M(1)$ is a super-martingale running backward with respect to the filtration $\mathcal{F}_{d+1} \subset \mathcal{F}_d \subset \dots \subset \mathcal{F}_1$. On one hand, if $k \in \mathcal{B}$, we have $V_+(k) = V_+(k-1)$, $V_-(k) = V_-(k-1)$ and thus $M(k) = M(k-1)$. On the other hand, if $k \in \mathcal{B}^c$, we have

$$M(k-1) = \frac{V_+(k) - I_k}{1 + V_-(k) - (1 - I_k)} = \frac{V_+(k) - I_k}{(V_-(k) + I_k) \vee 1},$$

where $I_k = I\{W_k > 0\}$ and $I\{\cdot\}$ is the indicator function. Let $d_0 = \#\{j : j \in \mathcal{B}^c\}$ with the inactive set $\mathcal{B}^c = \{j_1, j_2, \dots, j_{d_0}\}$. From Lemma 1 part (ii), we know that $I_{j_1}, I_{j_2}, \dots, I_{j_{d_0}}$ are iid Bernoulli(0.5) random variables. Thus conditioning on \mathcal{F}_k (hence $V_+(k)$, $V_-(k)$ are known), we have

$$\Pr(I_k = 1 | \mathcal{F}_k) = \Pr(I_k = 1 | V_+(k), V_-(k)) = \frac{V_+(k)}{V_+(k) + V_-(k)}.$$

Thus in the case $k \in \mathcal{B}^c$,

$$\begin{aligned} \mathbb{E}[M(k-1) | \mathcal{F}_k] &= \frac{V_+(k)}{V_+(k) + V_-(k)} \cdot \frac{V_+(k) - 1}{V_-(k) + 1} \\ &\quad + \frac{V_-(k)}{V_+(k) + V_-(k)} \cdot \frac{V_+(k)}{V_-(k) \vee 1}, \\ &= \begin{cases} \frac{V_+(k)}{V_-(k) + 1}, & \text{if } V_-(k) > 0, \\ V^+(k) - 1, & \text{if } V_-(k) = 0, \end{cases} \\ &= \begin{cases} M(k), & \text{if } V_-(k) > 0, \\ M(k) - 1, & \text{if } V_-(k) = 0. \end{cases} \end{aligned}$$

Therefore, $\mathbb{E}[M(k-1) | \mathcal{F}_k] \leq M(k)$, implying that $M(k), k = d+1, \dots, 1$ is a super-martingale with respect to $\{\mathcal{F}_k\}$. By definition, T_α is a stopping time with respect to the backward filtration $\{\mathcal{F}_k\}$. According to the optional stopping theorem for super-martingale, we know

$$\begin{aligned} \mathbb{E}[M(k_{T_\alpha})] &\leq \mathbb{E}[M(k_{d+1})] = \mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c, W_j > 0\}}{1 + \#\{j : j \in \mathcal{B}^c, W_j \leq 0\}} \right] \\ &= \mathbb{E} \left[\frac{X}{1 + d_0 - X} \right], \end{aligned}$$

where $X = \#\{j : j \in \mathcal{B}^c, W_j > 0\}$. Since $X \sim \text{Binomial}(d_0, 1/2)$, we have

$$\begin{aligned} \mathbb{E} \left[\frac{X}{1 + d_0 - X} \right] &= \sum_{k=1}^{d_0} \binom{d_0}{k} \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{d_0-k} \cdot \frac{k}{1 + d_0 - k} \\ &= \sum_{k=1}^{d_0} \binom{d_0}{k-1} \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{d_0-k} \\ &= \sum_{k=0}^{d_0-1} \binom{d_0}{k} \left(\frac{1}{2}\right)^{k+1} \left(\frac{1}{2}\right)^{d_0-k-1} \\ &\leq 1. \end{aligned}$$

Therefore, $\mathbb{E}[M(k_{T_\alpha})] \leq 1$. From (A.1), we have

$$\mathbb{E} \left[\frac{\#\{j : j \in \mathcal{B}^c \cap \widehat{\mathcal{A}}(T_\alpha)\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right] \leq \alpha \mathbb{E}[M(k_{T_\alpha})] \leq \alpha.$$

Since $\widehat{\mathcal{A}}(T_\alpha) \subset \widehat{\mathcal{A}}_1$, we have $\#\{j : j \in \mathcal{A}^c \cap \widehat{\mathcal{A}}(T_\alpha)\} = \#\{j : j \in \mathcal{B}^c \cap \widehat{\mathcal{A}}(T_\alpha)\}$. Then, we conclude

$$\mathbb{E} \left[\frac{\#\{j : j \in \mathcal{A}^c \cap \widehat{\mathcal{A}}(T_\alpha)\}}{\#\{j : j \in \widehat{\mathcal{A}}(T_\alpha)\} \vee 1} \right] \leq \alpha.$$

A.2. Proof of Theorem 5

Now we restrict ourselves to the subset $(\mathbf{X}^{(2)}, \mathbf{Y}^{(2)})$. For $j \in \widehat{\mathcal{A}}_1$, define $\widehat{\omega}_j = \widehat{\text{PC}}(\mathbf{X}_j^{(2)}, \mathbf{Y}^{(2)})^2$, $\tilde{\omega}_j = \widehat{\text{PC}}(\tilde{\mathbf{X}}_j^{(2)}, \mathbf{Y}^{(2)})^2$, $\omega_j = \text{PC}(X_j, \mathbf{y})^2$, and $\omega'_j = \text{PC}(\tilde{X}_j, \mathbf{y})^2$. Theorem 1 implies that

$$\begin{aligned} \Pr(|\widehat{\omega}_j - \omega_j| > c_3 n_2^{-\kappa}) &\leq 5c_1 \exp\{-c_4 n_2^{1-2\kappa}\} \text{ and} \\ \Pr(|\tilde{\omega}_j - \omega'_j| > c_3 n_2^{-\kappa}) &\leq 5c_1 \exp\{-c_4 n_2^{1-2\kappa}\}. \end{aligned}$$

Thus, we have

$$\begin{aligned} &\Pr(|\widehat{W}_j - W_j| > 2c_3 n_2^{-\kappa}) \\ &= \Pr(|\widehat{\omega}_j - \tilde{\omega}_j - (\omega_j - \omega'_j)| > 2c_3 n_2^{-\kappa}) \\ &\leq \Pr(|\widehat{\omega}_j - \omega_j| > c_3 n_2^{-\kappa}) + \Pr(|\tilde{\omega}_j - \omega'_j| > c_3 n_2^{-\kappa}) \\ &= O(\exp\{-c_4 n_2^{1-2\kappa}\}). \end{aligned} \quad (\text{A.2})$$

On the other hand, part (i) in Lemma 1 implies that $\omega_j = \omega'_j$ for all $j \in \mathcal{A}^c$. We have

$$\begin{aligned} &\Pr \left(\max_{j \in \mathcal{B}^c} |\widehat{W}_j| \leq 2c_3 n_2^{-\kappa} \right) \\ &= \Pr \left(\max_{j \in \mathcal{B}^c} |\widehat{\omega}_j - \omega_j + \omega'_j - \tilde{\omega}_j| \leq 2c_3 n_2^{-\kappa} \right) \\ &\geq \Pr \left(\max_{j \in \mathcal{B}^c} |\widehat{\omega}_j - \omega_j| \leq c_3 n_2^{-\kappa} \text{ and } \max_{j \in \mathcal{B}^c} |\tilde{\omega}_j - \omega_j| \leq c_3 n_2^{-\kappa} \right) \\ &\geq 1 - \sum_{j \in \mathcal{B}^c} \Pr(|\widehat{\omega}_j - \omega_j| \geq c_3 n_2^{-\kappa}) - \sum_{j \in \mathcal{B}^c} \Pr(|\tilde{\omega}_j - \omega_j| \geq c_3 n_2^{-\kappa}) \\ &\geq 1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\}), \end{aligned} \quad (\text{A.3})$$

where the last equality is implied by Theorem 1 and the assumption $n_2 > d$. Since $\min_{k \in \mathcal{B}} W_k \geq 4c_3 n_2^{-\kappa}$ and $d < n_2$, (A.2) implies that $\min_{k \in \mathcal{B}} \widehat{W}_k \geq 2c_4 n_2^{-\kappa}$ with probability at least $1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$. Together with (A.3), we know that $\min_{j \in \mathcal{B}} \widehat{W}_j > \max_{j \in \mathcal{B}^c} |\widehat{W}_j|$ with probability $1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$.

$O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$. In other words, with probability at least $1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$, the active features will be ranked ahead of the inactive features. Recall that to find the cutoff value T_α , we start from the smallest value $t = |\widehat{W}_{d+1}| = 0$, then move to the second smallest value $t = |\widehat{W}_d|$, then move to $t = |\widehat{W}_{d-1}|$, and so on. The procedure stops once it finds a value of t satisfying (13). Restrict on the event $\mathcal{E}' = \{\min_{j \in \mathcal{B}} \widehat{W}_j > \max_{j \in \mathcal{B}^c} |\widehat{W}_j|\}$, which holds with probability at least $1 - O(n_2 \exp\{-c_4 n_2^{1-2\kappa}\})$ and let $t_{\min} = \min_{j \in \mathcal{B}} |\widehat{W}_j| = \min_{j \in \mathcal{B}} \widehat{W}_j$. If $1/s \leq \alpha$, then

$$\frac{1 + \#\{j : \widehat{W}_j \leq -t_{\min}\}}{\#\{j : \widehat{W}_j \geq t_{\min}\}} = \frac{1 + 0}{s} \leq \alpha.$$

From the inequality above, we known this process must stop no later than t reaches t_{\min} and hence $T_\alpha \leq t_{\min}$. As a result,

$$\widehat{\mathcal{A}}(T_\alpha) = \{j : \widehat{W}_j \geq T_\alpha\} \supseteq \{j : \widehat{W}_j \geq t_{\min}\} = \mathcal{A}.$$

If $1/s > \alpha$, to satisfy (13), we must have $T_\alpha < t_{\min}$ or $T_\alpha > \max_j |\widehat{W}_j|$. $T_\alpha < t_{\min}$ indicates all active features are selected and $T_\alpha > \max_j |\widehat{W}_j|$ leads to an empty $\widehat{\mathcal{A}}(T_\alpha)$, as stated in the second part. Notice that $|\widehat{W}_{s+k}|$ satisfies the rule (13) is equivalent to that at most $\lfloor (k-1)/(s+1) \rfloor$ of the signs of $\{\widehat{W}_{s+1}, \dots, \widehat{W}_{s+k}\}$ are negative. Let a_k be the probability that $|\widehat{W}_{s+k}|$ satisfies (13) and b_k be the probability that the process stops at $|\widehat{W}_{s+k}|$ for $k = 1, \dots, d-s$. Let $a_0 = 0$, we have

$$a_k = \sum_{i=0}^{\lfloor (k-1)/(s+1) \rfloor} \binom{k}{i} \left(\frac{1}{2}\right)^k, b_k = a_k(1 - a_{k-1} - \dots - a_0).$$

It is easy to verify that for $s > 2$,

$$\begin{aligned} \Pr(\mathcal{A} \subset \widehat{\mathcal{A}}(T_\alpha) | \mathcal{E}') &= \sum_{k=1}^{d-s} b_k = \sum_{k=1}^{d-s} a_k(1 - a_{k-1} - \dots - a_0) \\ &\leq \sum_{k=1}^{\infty} a_k(1 - a_{k-1} - \dots - a_0) = C(s) < 1, \end{aligned}$$

where $0 < C(s) < 1$ is some constant only depending on s . This completes the proof.

Supplementary Materials

The supplementary pdf file provides additional technical proofs and numerical results. In Section S.1, we present the proofs of theoretical results in Section 2 as well as Theorem 6 in Section 3. In Section S.2, we report one additional simulated example and a real data application to complement the numerical analysis in the main document. The supplementary python file provides the codes of main functions for the article.

Acknowledgments

We would like to thank the AE and reviewers for their constructive comments, which lead to significant improvement of this work.

Funding

Liu's research was supported by National Natural Science Foundation of China (nos. 11771361, 11671334, 11871409), Basic Scientific Center Project 71988101 of National Science Foundation of China and JAS14007. Li's research was supported by NSF grants DMS 1820702 and 1953196.

References

Barber, R. F., and Candès, E. J. (2015), "Controlling the False Discovery Rate via Knockoffs," *The Annals of Statistics*, 43, 2055–2085. [5,6]

——— (2019), "A Knockoff Filter for High-Dimensional Selective Inference," *The Annals of Statistics*, 47, 2504–2537. [6,7,8]

Candès, E., Fan, Y., Janson, L., and Lv, J. (2018), "Panning for Gold: 'Model-X' Knockoffs for High-Dimensional Controlled Variable Selection," *Journal of the Royal Statistical Society, Series B*, 80, 551–577. [5]

Caruana, R. (1997), "Multitask Learning," *Machine Learning*, 28, 41–75. [2]

Chen, Z., Fan, J., and Li, R. (2018), "Error Variance Estimation in Ultrahigh-Dimensional Additive Models," *Journal of the American Statistical Association*, 113, 315–327. [6,11,12]

Fan, J., Feng, Y., and Song, R. (2011), "Nonparametric Independence Screening in Sparse Ultra-High-Dimensional Additive Models," *Journal of the American Statistical Association*, 106, 544–557. [1]

Fan, J., Ke, Y., Sun, Q., and Zhou, W.-X. (2019), "Farmtest: Factor-Adjusted Robust Multiple Testing With Approximate False Discovery Control," *Journal of the American Statistical Association*, 114, 1880–1893. [6]

Fan, J., Ke, Y., and Wang, K. (2020), "Factor-Adjusted Regularized Model Selection," *Journal of Econometrics*, 216, 71–85. [1]

Fan, J., Liao, Y., and Mincheva, M. (2013), "Large Covariance Estimation by Thresholding Principal Orthogonal Complements," *Journal of the Royal Statistical Society, Series B*, 75, 603–680. [5]

Fan, J., and Lv, J. (2008), "Sure Independence Screening for Ultrahigh-Dimensional Feature Space," *Journal of the Royal Statistical Society, Series B*, 70, 849–911. [1,8]

Fan, J., Samworth, R., and Wu, Y. (2009), "Ultrahigh Dimensional Feature Selection: Beyond the Linear Model," *Journal of Machine Learning Research*, 10, 2013–2038. [1]

Fan, J., and Song, R. (2010), "Sure Independence Screening in Generalized Linear Models With NP-Dimensionality," *The Annals of Statistics*, 38, 3567–3604. [1]

Fan, Y., Demirkaya, E., Li, G., and Lv, J. (2020), "Rank: Large-Scale Inference With Graphical Nonlinear Knockoffs," *Journal of the American Statistical Association*, 115, 362–379. [5,6,7]

Hall, P., and Miller, H. (2009), "Using Generalized Correlation to Effect Variable Selection in Very High Dimensional Problems," *Journal of Computational and Graphical Statistics*, 18, 533–550. [1]

He, X., Wang, L., and Hong, H. G. (2013), "Quantile-Adaptive Model-Free Variable Screening for High-Dimensional Heterogeneous Data," *The Annals of Statistics*, 41, 342–369. [1]

Hoeffding, W. (1948), "A Non-Parametric Test of Independence," *The Annals of Mathematical Statistics*, 19, 546–557. [3]

Ke, Y., Minsker, S., Ren, Z., Sun, Q., and Zhou, W.-X. (2019), "User-Friendly Covariance Estimation for Heavy-Tailed Distributions," *Statistical Science*, 34, 454–471. [5]

Li, R., Zhong, W., and Zhu, L. (2012), "Feature Screening via Distance Correlation Learning," *Journal of the American Statistical Association*, 107, 1129–1139. [1,2,3,4,8,9]

Liu, J., Li, R., and Wu, R. (2014), "Feature Selection for Varying Coefficient Models With Ultrahigh-Dimensional Covariates," *Journal of the American Statistical Association*, 109, 266–274. [1]

Liu, J., Zhong, W., and Li, R. (2015), "A Selective Overview of Feature Screening for Ultrahigh-Dimensional Data," *Science China Mathematics*, 58, 1–22. [1]

Liu, W., and Li, R. (2020), "Variable Selection and Feature Screening," in *Macroeconomic Forecasting in the Era of Big Data* (Vol. 52), ed. P. Fuleky, Cham: Springer, pp. 293–326. [1]

Lv, J., and Liu, J. S. (2014), "Model Selection Principles in Misspecified Models," *Journal of the Royal Statistical Society, Series B*, 76, 141–167. [1]

Ma, S., Li, R., and Tsai, C.-L. (2017), "Variable Screening via Quantile Partial Correlation," *Journal of the American Statistical Association*, 112, 650–663. [1]

Mai, Q., and Zou, H. (2012), "The Kolmogorov Filter for Variable Screening in High-Dimensional Binary Classification," *Biometrika*, 100, 229–234. [2]

——— (2015), "The Fused Kolmogorov Filter: A Nonparametric Model-Free Screening Method," *The Annals of Statistics*, 43, 1471–1497. [1]

Pan, R., Wang, H., and Li, R. (2016), "Ultrahigh-Dimensional Multiclass Linear Discriminant Analysis by Pairwise Sure Independence Screening," *Journal of the American Statistical Association*, 111, 169–179. [2]

Song, R., Lu, W., Ma, S., and Jessie Jeng, X. (2014), "Censored Rank Independence Screening for High-Dimensional Survival Data," *Biometrika*, 101, 799–814. [2,4]

Székely, G. J., and Rizzo, M. L. (2010), "Brownian Distance Covariance," *The Annals of Applied Statistics*, 3, 1236–1265. [3]

——— (2014), "Partial Distance Correlation With Methods for Dissimilarities," *The Annals of Statistics*, 42, 2382–2412. [8]

Wang, H. (2009), "Forward Regression for Ultra-High Dimensional Variable Screening," *Journal of the American Statistical Association*, 104, 1512–1524. [11]

——— (2012), "Factor Profiled Sure Independence Screening," *Biometrika*, 99, 15–28. [1]

Wu, Y., and Yin, G. (2015), "Conditional Quantile Screening in Ultrahigh-Dimensional Heterogeneous Data," *Biometrika*, 102, 65–76. [1]

Xue, L., and Zou, H. (2011), "Sure Independence Screening and Compressed Random Sensing," *Biometrika*, 98, 371–380. [1]

Zhao, S. D., and Li, Y. (2012), "Principled Sure Independence Screening for Cox Models With Ultra-High-Dimensional Covariates," *Journal of Multivariate Analysis*, 105, 397–411. [2,4]

Zhu, L.-P., Li, L., Li, R., and Zhu, L.-X. (2011), "Model-Free Feature Screening for Ultrahigh-Dimensional Data," *Journal of the American Statistical Association*, 106, 1464–1475. [1]

Zhu, L., Xu, K., Li, R., and Zhong, W. (2017), "Projection Correlation Between Two Random Vectors," *Biometrika*, 104, 829–843. [2,3,4,13]