

Optimal and Safe Estimation for High-Dimensional Semi-Supervised Learning

Siyi Deng^a, Yang Ning^a, Jiwei Zhao^b, and Heping Zhang^c

^aDepartment of Statistics and Data Science, Cornell University, Ithaca, NY; ^bDepartment of Biostatistics and Medical Informatics, University of Wisconsin-Madison, Madison, WI; ^cDepartment of Biostatistics, Yale University, New Haven, CT

ABSTRACT

We consider the estimation problem in high-dimensional semi-supervised learning. Our goal is to investigate when and how the unlabeled data can be exploited to improve the estimation of the regression parameters of linear model in light of the fact that such linear models may be misspecified in data analysis. We first establish the minimax lower bound for parameter estimation in the semi-supervised setting, and show that this lower bound cannot be achieved by supervised estimators using the labeled data only. We propose an optimal semi-supervised estimator that can attain this lower bound and therefore improves the supervised estimators, provided that the conditional mean function can be consistently estimated with a proper rate. We further propose a safe semi-supervised estimator. We view it safe, because this estimator is always at least as good as the supervised estimators. We also extend our idea to the aggregation of multiple semi-supervised estimators caused by different misspecifications of the conditional mean function. Extensive numerical simulations and a real data analysis are conducted to illustrate our theoretical results. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received July 2021
Accepted October 2023

KEYWORDS

High dimensionality; Model aggregation; Model misspecification; Optimal estimation; Safe estimation; Semi-supervised learning

1. Introduction

Semi-supervised learning is an emerging research area in statistics and machine learning (Zhu 2005; Chapelle, Scholkopf, and Zien 2006), and can have a great potential in electronic health records (EHR) based studies for clinical research. In these types of studies, one major challenge is the lack of gold-standard health outcomes or phenotypes (Kohane 2011). The validated phenotypes are often obtained by manual chart reviews that are prohibitively expensive (Liao et al. 2010); therefore, only can a very small subset of patients be annotated by experts in reality. For the rest of the patients, their covariate information, often high-dimensional (Weisenthal et al. 2018; Gensheimer et al. 2019; Abdullah et al. 2020), is only available.

Developing efficient statistical methods to analyze such data is a timely and important problem. Let Y denote the outcome variable and X the p -dimensional covariates. In addition to n iid samples $(Y_1, X_1), \dots, (Y_n, X_n) \sim (Y, X)$, we also observe N iid data consisting of only covariates, $X_{n+1}, \dots, X_{N+n} \sim X$. Following the convention, the former is referred to as labeled data and the latter is called unlabeled data. We also denote $\mathbf{Y} = (Y_1, \dots, Y_n)^T \in \mathbb{R}^n$ and $\mathbf{X} = (X_1, \dots, X_n)^T \in \mathbb{R}^{n \times p}$ the outcomes and covariates from the labeled data, and $\tilde{\mathbf{X}} = (X_1, \dots, X_{N+n})^T \in \mathbb{R}^{(N+n) \times p}$ the covariates from both. In this work, we focus on high-dimensional regression problems; namely, p can be much larger than n . The size of the unlabeled data N is allowed but not required to be larger than n .

We consider the so-called assumption lean regression framework (Buja et al. 2019; Berk et al. 2019),

$$Y = f(X) + \epsilon, \quad (1.1)$$

where $f(X) = E(Y|X)$ is the unknown conditional mean function, ϵ is the random error independent of $X \in \mathbb{R}^p$ with $\mathbb{E}(\epsilon) = 0$, $\mathbb{E}(\epsilon^2) = \sigma^2$, and σ^2 is an unknown parameter. We consider the random design and assume that X and Y are centered with $\mathbb{E}(X) = 0$ and $\mathbb{E}(f(X)) = 0$. On one hand, we would like to put as fewer assumptions as possible on $f(X)$ to enable model flexibility. On the other hand, for the purpose of interpretability, we often fit simple parametric models such as linear regression to explain the association between Y and X . To meet both ends, we consider linear regression as a working model where the true data generating process follows (1.1). Since $\mathbb{E}[(Y - X^T \theta)^2] = \mathbb{E}[(f(X) - X^T \theta)^2] + \sigma^2$, the regression coefficients in a linear model correspond to the $L_2(\mathbb{P})$ projection of $f(X)$ onto the linear space spanned by X , that is,

$$\theta^* = \arg \min_{\theta \in \mathbb{R}^p} \mathbb{E}[(f(X) - X^T \theta)^2]$$

that delineates the linear dependence between Y and X . We do not include intercept in θ^* simply because $\mathbb{E}(X) = 0$ and $\mathbb{E}(f(X)) = 0$. Our goal here is to estimate the high dimensional parameter θ^* .

In the supervised setting with n labeled observations, a number of penalized estimators have been proposed to estimate θ^* , such as lasso (Tibshirani 1996) and Dantzig selector (Candes and Tao 2007). While significant progress has been made toward understanding the estimation in the fully supervised setting, there is very limited research in the semi-supervised setting. It is

important to observe that under (1.1), since the linear regression is the working model, the covariate X is no longer the ancillary statistic for the regression parameter θ^* . Therefore, the covariate X in the unlabeled data, usually with a much bigger sample size than the labeled data, is informative and may be beneficial for estimating θ^* .

Our first contribution is to establish the minimax lower bound for estimating θ^* in the semi-supervised setting. In particular, to derive this lower bound, we assume $f(\cdot)$ is unknown but belongs to some specific model class, such that methods for estimating $f(\cdot)$ are available in the existing literature. Based on the lower bound, when N is sufficiently large, the error term due to the model misspecification or equivalently the nonlinearity of $f(\cdot)$ becomes negligible. This reveals one potential benefit from using the unlabeled data in that the estimation of θ^* can be more robust to the model misspecification. Moreover, we show that the fully supervised estimators (e.g., lasso and Dantzig selector) do not attain this lower bound. Our second contribution is to propose a new semi-supervised estimator that achieves this lower bound under some conditions. In [Theorem 3.2](#), we show that the rate of our estimator depends on whether the unknown $f(\cdot)$ can be consistently estimated. When $f(\cdot)$ belongs to some specific model class so that it can be consistently estimated with a proper rate, the proposed estimator attains the minimax lower bound up to some logarithmic factor, and therefore improves the rate of the supervised estimators. When $f(\cdot)$ is misspecified, however, the rate of our semi-supervised estimator becomes sub-optimal and may be even slower than the supervised estimators. Our third contribution is to develop a general two-step refitting procedure that further improves the above semi-supervised estimator. The resulting estimator is guaranteed to be no worse than the supervised estimators no matter $f(\cdot)$ is misspecified or not, and remains minimax rate-optimal when $f(\cdot)$ belongs to some specific model class so that can be consistently estimated with a proper rate, hence, it provides a safe use of the unlabeled data. Therefore, we call it the safe semi-supervised estimator. In empirical studies one may encounter the situation that, while all misspecified, different estimates of $f(\cdot)$ are available. We further extend the idea of creating safe semi-supervised estimator to the aggregation of multiple semi-supervised estimators under such a situation. The convergence rate of the aggregated estimator is guaranteed to be no worse than any of the un-aggregated semi-supervised estimators. Overall, our goal is to exploit both safe and efficient use of the unlabeled data by developing semi-supervised estimators of θ^* such that their convergence rates are faster, or at least no slower, than the standard supervised estimators (e.g., lasso and Dantzig selector).

1.1. Related Work

In computer science, a large number of classification algorithms have been developed under semi-supervised setting, which mainly focus on data with discrete labels; see [Zhu \(2005\)](#) and [Chapelle, Scholkopf, and Zien \(2009\)](#) for some surveys. Common assumptions such as manifold assumption and cluster assumption were made in the literature in order to obtain fast rate of convergence in classification ([Rigollet 2006](#)). In nonparametric regression problem, [Wasserman and Lafferty \(2008\)](#) showed that unlabeled data do not always help to improve the rate of the mean squared error, but with semi-supervised smoothness assumption the estimator with faster rate could be developed.

([2008](#)) showed that unlabeled data do not always help to improve the rate of the mean squared error, but with semi-supervised smoothness assumption the estimator with faster rate could be developed.

More recently, [Van Engelen and Hoos \(2020\)](#) and [Yang et al. \(2021\)](#) provided comprehensive surveys on many up-to-date developments in semi-supervised learning, especially with deep neural networks. For example, the methods named as Mixmatch ([Berthelot et al. 2019](#)) and Fixmatch ([Sohn et al. 2020](#)) leverage unlabeled data through pseudo-labeling and consistency regularization to assist the prediction model. The self-training method, which trains a model to fit pseudo-labels predicted by previously learned models, has also been developed for semi-supervised learning ([Xie et al. 2020](#); [Chen et al. 2020](#); [Wei et al. 2020](#)). The performance of many modern semi-supervised learning algorithms in some image classification tasks is compared and discussed by [Oliver et al. \(2018\)](#).

There are also progress considering how to make use of the unlabeled data to obtain an estimator with a smaller asymptotic variance, when the dimension p is fixed and small. For example, [Zhang, Brown, and Cai \(2019\)](#), [Azriel et al. \(2016\)](#), and [Chakrabortty and Cai \(2018\)](#) investigated how to incorporate the unlabeled data to improve the estimation efficiency for the population mean $\mathbb{E}(Y)$ and regression coefficients in a working linear regression.

With high-dimensional data, [Alquier and Hebiri \(2012\)](#) proposed a transductive version of lasso and Dantzig selector in the semi-supervised setting. They showed that the transductive estimators have the same rate as the supervised ones. More recently, [Bellec et al. \(2018\)](#) proposed a modified lasso estimator and showed that the excess risk of their estimator in prediction has the same rate of convergence as the supervised lasso estimator. These existing theoretical results neither confirm nor deny the existence of estimators with improved statistical rate when unlabeled data are available. To the best of our knowledge, it remains an open problem of when and how one can develop a semi-supervised estimator with improved statistical rate by exploiting the available unlabeled data. We bridge this gap by showing the minimax lower bound and proposing new semi-supervised estimators.

1.2. Organization of the Article

The rest of this article is organized as follows. In [Section 2](#), we give the minimax lower bound for semi-supervised estimation. In [Section 3](#), we introduce the optimal semi-supervised estimator and its corresponding upper bound. In [Section 4](#) we devote ourselves into the safe semi-supervised estimator, where we propose a two-step procedure regardless of the quality of the estimated conditional mean model. In [Section 5](#) we extend the idea to aggregation of multiple semi-supervised estimators caused by different misspecifications of the conditional mean model. Numerical experiments and a real data application are in [Sections 6](#) and [7](#), respectively. All the technical proofs are contained in the supplementary materials.

1.3. Notations

Let $P_{X,Y}$ and P_X denote the joint distribution of (X, Y) and the marginal distribution of X , respectively. For $v = (v_1, \dots, v_p)^T \in$

\mathbb{R}^p , we define $\|v\|_0 = |\text{supp}(v)|$ where $\text{supp}(v) = \{i : v_i \neq 0\}$ and $|A|$ is the cardinality of a set A , $\|v\|_q = (\sum_{i=1}^p |v_i|^q)^{1/q}$ for $1 \leq q < \infty$, and $\|v\|_\infty = \max_{1 \leq i \leq p} |v_i|$. Denote $v^{\otimes 2} = vv^T$. For a matrix $\mathbf{M} = [M_{ij}]$, \mathbf{M}_i and \mathbf{M}_j denote the i th row and j th column, respectively. Define $\|\mathbf{M}\|_{\max} = \max_{ij} |M_{ij}|$, $\|\mathbf{M}\|_1 = \max_j \sum_i |M_{ij}|$, $\|\mathbf{M}\|_\infty = \max_i \sum_j |M_{ij}|$. If the matrix \mathbf{M} is symmetric, then $\Lambda_{\min}(\mathbf{M})$ and $\Lambda_{\max}(\mathbf{M})$ are the minimal and maximal eigenvalues of \mathbf{M} . We denote \mathbf{I}_p the $p \times p$ identity matrix. For $S \subseteq \{1, \dots, p\}$, let $v_S = \{v_k : k \in S\}$ and S^c be the complement of S . For matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ and index set $D \subseteq \{1, \dots, n\}$, $\mathbf{X}_D = \{X_i : i \in D\}^T \in \mathbb{R}^{|D| \times p}$. For a function f , let $\|f\|_2 = \sqrt{\mathbb{E}[f(X)^2]}$ denote the $L_2(\mathbb{P})$ norm of f .

For two positive sequences a_n and b_n , we write $a_n \asymp b_n$ if $C \leq a_n/b_n \leq C'$ for some $C, C' > 0$. Similarly, we use $a \lesssim b$ to denote $a \leq Cb$ for some constant $C > 0$. Given $a, b \in \mathbb{R}$, let $a \vee b$ and $a \wedge b$ denote the maximum and minimum of a and b .

2. Minimax Lower Bound for Semi-Supervised Estimation

The semi-supervised learning setting refers to that we observe n iid copies of (Y, X) and additional N iid copies of X , where the distributions of X in both labeled and unlabeled data are the same. In [Theorem 3.2](#) presented in [Section 3](#), we will rigorously show that the unknown conditional mean function $f(\cdot)$ plays an important role in assessing the optimality of semi-supervised estimators. However, the correct specification and consistent estimation of $f(\cdot)$ under high dimensionality is by no means a trivial problem. In the literature, methods for consistently estimating $f(\cdot)$ with a proper rate are only available when $f(\cdot)$ belongs to some specific model class. In that regard, when analyzing the minimax lower bound here, we assume that $f(\cdot)$ belongs to one of the following two model classes.

We first define the pairwise interaction model $\mathcal{F}_{\text{pairwise}} = \{\sum_{j=1}^p \gamma_j X_j + \sum_{1 \leq j \leq k \leq p} \gamma_{jk} X_j X_k\}$, where $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p, \gamma_{11}, \gamma_{12}, \dots, \gamma_{pp}) \in \mathbb{R}^{p+p(p+1)/2}$ are unknown parameters satisfying $\|\boldsymbol{\gamma}\|_0 \leq \bar{s}$. In practice, the conditional mean function $f(X)$ is usually nonlinear in X . In $\mathcal{F}_{\text{pairwise}}$, we account for the nonlinearity by incorporating the quadratic terms and the pairwise interactions. To mitigate the model complexity, the parameter $\boldsymbol{\gamma}$ is assumed to be \bar{s} -sparse ([Zhao and Leng 2016](#)). We refer to [Supplement S.4](#) for further discussions. Thus, we define the class of joint distributions of (X, Y) as

$$\begin{aligned} \mathcal{P}_{\Phi, \sigma}^{\text{pairwise}} &= \{P_{X,Y} \mid Y = f(X) + \epsilon, f(\cdot) \in \mathcal{F}_{\text{pairwise}}, \\ &\quad \|\boldsymbol{\theta}^*\|_0 \leq s, \text{ var}(\epsilon) = \sigma^2, \\ &\quad \mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2 \leq \Phi^2, \text{ and } P_X \in \mathcal{P}_X\}, \end{aligned}$$

where $\boldsymbol{\theta}^*$ implicitly depends on the distribution $P_{X,Y}$, the parameter s controls the sparsity of $\boldsymbol{\theta}^*$, $\mathcal{P}_X = \{P_X \mid \mathbb{E}[X] = 0, \text{ var}(X_j) = 1 \text{ and } \Lambda_{\min}(\text{cov}(X)) \geq C_{\min} > 0\}$ with some constant C_{\min} . For notational simplicity, we write $\mathbb{E}(\cdot)$ for $\mathbb{E}_{P_{X,Y}}(\cdot)$. We note that, $\mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}$ is indexed by two non-negative parameters Φ^2 and σ^2 , where the former controls the magnitude of model misspecification $f(X) - X^T \boldsymbol{\theta}^*$ or equivalently the nonlinearity of $f(X)$ in the second moment and the latter is the variance of ϵ . In particular, we allow Φ^2 to grow with n in our framework.

The second model class we consider is the additive model $\mathcal{F}_{\text{additive}} = \{\sum_{j=1}^p f_j(X_j)\}$, where f_j 's are unknown second-order-smooth functions ([Lin and Zhang 2006](#); [Meier, Van de Geer, and Bühlmann 2009](#); [Huang, Horowitz, and Wei 2010](#); [Raskutti, Wainwright, and Yu 2012](#)). To ease the presentation, we defer the definition of ℓ -smooth functions to [Supplement S.4](#). Similarly, we assume the number of nonzero functions is bounded by \bar{s} . Compared to $\mathcal{F}_{\text{pairwise}}$, the additive model does not allow interactions among covariates but offers more flexibility in associating each component X_j with Y . Similar to $\mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}$, we can define the class of distributions $\mathcal{P}_{\Phi, \sigma}^{\text{additive}}$, where we replace $f \in \mathcal{F}_{\text{pairwise}}$ with $f \in \mathcal{F}_{\text{additive}}$ in the definition.

The following theorem offers the lower bound for the convergence rate of any estimator of $\boldsymbol{\theta}^*$ over the classes of distributions $\mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}$ or $\mathcal{P}_{\Phi, \sigma}^{\text{additive}}$, in the semi-supervised setting. Throughout the paper without causing confusion, we use C, C', c_1 , and c_2 , etc. to denote generic constants whose values can change from time to time.

Assumption 2.1. Assume that $s \log(p/s) \leq Cn$ for some constant C , $4 \leq s \leq (n-1)/4$, and the sparsity level in $\mathcal{F}_{\text{pairwise}}$ and $\mathcal{F}_{\text{additive}}$ satisfies $\bar{s} \geq s$.

Theorem 2.2. Under [Assumption 2.1](#), we have that for any $1 \leq q \leq \infty$,

$$\inf_{\widehat{\boldsymbol{\theta}}} \sup_{P_{X,Y} \in \mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}} \mathbb{P}_{P_{X,Y}} \left[\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_q \geq c_1 s^{1/q} \left(\Phi \sqrt{\frac{\log(p/s)}{n+N}} + \sigma \sqrt{\frac{\log(p/s)}{n}} \right) \right] > c_2, \quad (2.1)$$

where $\inf_{\widehat{\boldsymbol{\theta}}}$ denotes the infimum over all estimators based on the labeled data (Y_i, X_i) for $1 \leq i \leq n$ and unlabeled data X_i for $n+1 \leq i \leq n+N$, and c_1 and c_2 are some positive constants. Here we denote $s^{1/\infty} = 1$. In addition, the same lower bound (2.1) holds when we replace $P_{X,Y} \in \mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}$ with $P_{X,Y} \in \mathcal{P}_{\Phi, \sigma}^{\text{additive}}$.

Remark 2.3. The lower bound (2.1) in [Theorem 2.2](#) is obtained by restricting $f(\cdot)$ to be in either $\mathcal{F}_{\text{pairwise}}$ or $\mathcal{F}_{\text{additive}}$. Indeed, this is a stronger result than the case without such a restriction. [Theorem 2.2](#) implies, if one considers the class of joint distributions of (X, Y) as

$$\begin{aligned} \mathcal{P}_{\Phi, \sigma} &= \{P_{X,Y} \mid Y = f(X) + \epsilon, \|f\|_2 < \infty, \|\boldsymbol{\theta}^*\|_0 \leq s, \\ &\quad \text{var}(\epsilon) = \sigma^2, \mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2 \leq \Phi^2, \text{ and } P_X \in \mathcal{P}_X\} \end{aligned}$$

without specifying the model class of $f(\cdot)$, the minimax lower bound remains the same as in (2.1). This lower bound consists of two components. Up to some absolute constants, the first term $s^{1/q} \Phi \sqrt{\log(p/s)/(n+N)}$ corresponds to the error due to potential model misspecification and the second term $s^{1/q} \sigma \sqrt{\log(p/s)/n}$ comes from the uncertainty inherited from the randomness of the error ϵ , which always exists even if the regression function is linear $f(X) = X^T \boldsymbol{\theta}^*$. In this case, we have $\Phi = 0$ and the lower bound agrees with the existing result for sparse linear regression ([Verzelen 2012](#); [Bellec, Lecué, and Tsybakov 2018](#)).

Remark 2.4. The sample size of the unlabeled data N plays an important role in the lower bound (2.1). In Supplement S.4.1, we show that under some conditions, $\Phi^2 \asymp s$ so $\Phi \rightarrow \infty$ as the sparsity grows and $\frac{\Phi}{\sigma} \sqrt{\frac{n}{n+N}} \rightarrow \infty$ may happen. In this case, the dominating term in the lower bound $s^{1/q} \Phi \sqrt{\log(p/s)/(n+N)}$ can be reduced as N increases. If N is sufficiently large such that $\frac{\Phi}{\sigma} \sqrt{\frac{n}{n+N}} \rightarrow c < \infty$, the lower bound attains its minimum $s^{1/q} \sigma \sqrt{\log(p/s)/n}$, which can be viewed as the irreducible error in the semi-supervised setting since a further increase of N would no longer decrease the lower bound. As an illustration, we plot the lower bound in Figure 1 of Supplement S.2.

Before delving into our proposed estimators, we briefly summarize some known properties of supervised estimators which shall be useful later. The supervised Dantzig selector is defined as

$$\hat{\boldsymbol{\theta}}_D = \arg \min \|\boldsymbol{\theta}\|_1, \text{ s.t. } \left\| \frac{1}{n} \sum_{i=1}^n (Y_i - X_i^T \boldsymbol{\theta}) X_i \right\|_\infty \leq \lambda_D, \quad (2.2)$$

where λ_D is a tuning parameter. It is shown in Lemma S.4 of Supplement S.3 that with high probability

$$\|\hat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*\|_1 = O_p \left\{ s(\Phi + \sigma) \sqrt{\frac{\log p}{n}} \right\}. \quad (2.3)$$

Under the condition $N \gg n$, the lower bound in (2.1) is strictly smaller in order than the upper bound (2.3) if and only if $\Phi/\sigma \rightarrow \infty$. In this case, the supervised estimator $\hat{\boldsymbol{\theta}}_D$ does not attain the lower bound and is thus sub-optimal in the minimax sense; see Figure 1 in Supplement S.2. Similarly, the supervised lasso estimator is defined as

$$\hat{\boldsymbol{\theta}}_L = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} \frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T \boldsymbol{\theta})^2 + \lambda_L \|\boldsymbol{\theta}\|_1, \quad (2.4)$$

where λ_L is a tuning parameter. The same upper bound as in (2.3) can be derived similarly. The Dantzig selector and the lasso estimator are theoretically equivalent (Bickel, Ritov, and Tsybakov 2009). To better illustrate the motivation of how to incorporate the unlabeled data, we next introduce our semi-supervised estimator with the formulation of the Dantzig selector in Section 3.

3. Optimal Semi-Supervised Estimator

3.1. Motivation and the Key Step

To motivate our estimator, we first briefly explain how the convergence rate of $\hat{\boldsymbol{\theta}}_D$ in (2.2) is derived. Following the standard argument in Bickel, Ritov, and Tsybakov (2009), the Dantzig selector satisfies $\|\hat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*\|_1 = O_p(s\lambda_D)$, where the tuning parameter $\lambda_D \gtrsim \|\frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*)\|_\infty$. In the proof of Lemma S.4, we further show that $\|\frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*)\|_\infty \lesssim \sqrt{\frac{\log p}{n}} \{\mathbb{E}(Y_i - X_i^T \boldsymbol{\theta}^*)^2\}^{1/2}$ with high probability. The desired bound (2.3) is obtained by noting that

$$\mathbb{E}(Y_i - X_i^T \boldsymbol{\theta}^*)^2 = \mathbb{E}(Y_i - f(X_i))^2 + \mathbb{E}(f(X_i) - X_i^T \boldsymbol{\theta}^*)^2 \leq \sigma^2 + \Phi^2. \quad (3.1)$$

In view of (2.3) and Remark 2.3, we see that the slow rate of $\hat{\boldsymbol{\theta}}_D$ is driven by the sup-norm of the score function $\|\frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*)\|_\infty$.

To find an estimator with the improved rate, our key idea is to construct a modified score function. To this end, we decompose the score function of $\hat{\boldsymbol{\theta}}_D$ as

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*) &= \frac{1}{n} \sum_{i=1}^n X_i(Y_i - f(X_i)) \\ &\quad + \frac{1}{n} \sum_{i=1}^n X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*). \end{aligned}$$

We propose to replace the last term as $\frac{1}{n+N} \sum_{i=1}^{n+N} X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*)$, the sample average over both labeled and unlabeled data. Apparently, it is a consistent estimator of $\mathbb{E}[X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*)] = -\frac{1}{2} \mathbb{E}[\frac{\partial}{\partial \boldsymbol{\theta}}(f(X_i) - X_i^T \boldsymbol{\theta}^*)^2]$ with a faster rate. Thus, the unlabeled data can help estimate the expectation of the gradient of the model misspecification error. This explains why the unlabeled data may help in the case of misspecification. This leads to the following modified score function

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n X_i(Y_i - f(X_i)) + \frac{1}{n+N} \sum_{i=1}^{n+N} X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*) \\ = \bar{\boldsymbol{\xi}} - \hat{\boldsymbol{\Sigma}}_{n+N} \boldsymbol{\theta}^*, \end{aligned} \quad (3.2)$$

where $\hat{\boldsymbol{\Sigma}}_{n+N} = \frac{1}{n+N} \sum_{i=1}^{n+N} X_i^{\otimes 2}$ and

$$\bar{\boldsymbol{\xi}} = \frac{1}{n} \sum_{i=1}^n X_i Y_i - \frac{1}{n} \sum_{i=1}^n X_i f(X_i) + \frac{1}{n+N} \sum_{i=1}^{n+N} X_i f(X_i). \quad (3.3)$$

Further insight on the modified score function can be found in Supplement S.4.5.

3.2. Computation of $\bar{\boldsymbol{\xi}}$

To compute $\bar{\boldsymbol{\xi}}$, we need to find an estimator for $f(\cdot)$, the unknown conditional mean function. In the rest of the article, we use $\hat{h}(\cdot)$ to denote the estimate of the conditional mean function and will discuss some examples in Remark 3.3. To account for the possible model misspecification of the unknown conditional mean function, we assume that there exists a function $h(\cdot)$ with $\|h\|_2 < \infty$ such that the estimate $\hat{h}(\cdot)$ converges to $h(\cdot)$ in the $L_2(\mathbb{P})$ norm. We refer to h as a conditional mean model. When the conditional mean model is correctly specified, we would expect that $f = h$ and $\hat{h}(\cdot)$ is consistent for $f(\cdot)$.

A serious challenge may arise from deriving the theoretical property of our proposed semi-supervised estimator if we use all data to obtain $\hat{h}(\cdot)$ due to the dependence between the estimator $\hat{h}(\cdot)$ and the data (X_i, Y_i) in the sample average from $\bar{\boldsymbol{\xi}}$. To bypass this challenge, we adopt the cross-fitting technique that was devised for semiparametric estimation problems (Bickel 1982; Schick 1986) as well as for high-dimensional data (Robins et al. 2017; Chernozhukov et al. 2018). For notational simplicity, we denote by D^* the labeled data and D the full dataset. Without loss of generality, we split the labeled data D^* into two folds D_1^* and D_2^* with size $n_1 = n_2 = n/2$. Similarly, we split the unlabeled data into two folds U_1 and U_2 with size $N_1 = N_2 = N/2$.

Merging U_1 and U_2 with D_1^* and D_2^* , respectively, we obtain two independent datasets $D_1 = D_1^* \cup U_1$ and $D_2 = D_2^* \cup U_2$. Next, for $j = \{1, 2\}$, we train the estimator \widehat{h}^{-j} using the data $D^* \setminus D_j^*$ and then construct

$$\widehat{\xi}_j = \frac{1}{n_j} \sum_{i \in D_j^*} X_i Y_i - \frac{1}{n_j} \sum_{i \in D_j^*} X_i \widehat{h}^{-j}(X_i) + \frac{1}{n_j + N_j} \sum_{i \in D_j} X_i \widehat{h}^{-j}(X_i). \quad (3.4)$$

In view of the modified score function (3.2), replacing $\widehat{\xi}$ with $\widehat{\xi} = (\widehat{\xi}_1 + \widehat{\xi}_2)/2$, we propose the following semi-supervised Dantzig selector

$$\widehat{\theta}_{SD} = \arg \min \|\theta\|_1, \text{ s.t. } \|\widehat{\Sigma}_{n+N}\theta - \widehat{\xi}\|_\infty \leq \lambda_{SD}. \quad (3.5)$$

Similarly, we define the semi-supervised lasso estimator as

$$\widehat{\theta}_{SL} = \arg \min_{\theta \in \mathbb{R}^p} \theta^T \widehat{\Sigma}_{n+N} \theta - 2\widehat{\xi}^T \theta + 2\lambda_{SL} \|\theta\|_1. \quad (3.6)$$

3.3. Theoretical Property of the Proposed Estimator

We develop the theoretical property for the proposed optimal semi-supervised estimator.

Assumption 3.1. We make the following assumptions:

- (A1) $\Sigma^{-1/2}X$ is a zero mean sub-Gaussian vector with bounded sub-Gaussian norm and $\text{cov}(X) = \Sigma$ has smallest eigenvalue $\Lambda_{\min}(\Sigma) \geq C_{\min} > 0$ for some positive constant C_{\min} . Moreover, $\max_{1 \leq j \leq p} \Sigma_{jj} = O(1)$.
- (A2) $\max_{1 \leq i \leq n+N} \|X_i\|_\infty \leq K_1$ where we allow K_1 to diverge with (n, N, p) .
- (A3) $\mathbb{E}(\epsilon^2) = \sigma^2$ and $\mathbb{E}[(f(X) - X^T\theta^*)^2] \leq \Phi^2$.
- (A4) θ^* is s -sparse with $\|\theta^*\|_0 \leq s$, and $\frac{s \log p}{n+N} = O(1)$.

Assumption (A1) is a standard technical condition for X in order to verify the restricted eigenvalue (RE) condition (Bickel, Ritov, and Tsybakov 2009). Assumption (A2) imposes the boundedness of the covariates, which simplifies the analysis when the linear model is misspecified (Bühlmann and Van de Geer 2015). In particular, when X_i is uniformly bounded, K_1 becomes a constant. If each component of X_i is Gaussian or sub-Gaussian, Assumption (A2) still holds with high probability with $K_1 = C\sqrt{\log[p(n+N)]}$ for some constant C . Assumption (A3) only requires the existence of the second moment of ϵ and $f(X) - X^T\theta^*$. We note that, unlike Bickel, Ritov, and Tsybakov (2009), we do not assume the residual $Y - X^T\theta^*$ to be sub-Gaussian. This is because the residual in the misspecified model $Y - X^T\theta^* = \epsilon + (f(X) - X^T\theta^*)$ contains the nonlinear term $f(X) - X^T\theta^*$ which can be large. While we only assume the moment condition in Assumption (A3), the boundedness in Assumption (A2) enables us to apply the Nemirovski moment inequality (Lemma S.2) to control the deviation of the sample estimates from their population. Assumption (A4) is the sparsity condition. In particular, Bühlmann and Van de Geer (2015) provided some sufficient conditions on $f(X)$ and the distribution of X under which θ^* is sparse in the misspecified model. We further require $\frac{s \log p}{n+N} = O(1)$ to verify the RE condition under the random design; see Lemma S.3.

Given Assumption 3.1, we establish the convergence rate of the semi-supervised Dantzig selector $\widehat{\theta}_{SD}$ in (3.5). By Bickel,

Ritov, and Tsybakov (2009), one can easily show that the same error bounds hold for the semi-supervised lasso estimator $\widehat{\theta}_{SL}$. For simplicity, we only present the asymptotic results for $\widehat{\theta}_{SD}$, where $n, p \rightarrow \infty$ and N can be either fixed or tends to infinity as well.

Theorem 3.2. Suppose Assumption 3.1 holds and the estimator $\widehat{h}^{-j}(\cdot)$ satisfies

$$\|\widehat{h}^{-j} - h\|_2 = O_p(b_n),$$

for $j = 1, 2$, where b_n is a deterministic sequence. Denote $G_h = \|h - f\|_2$. With some tuning parameter $\lambda_{SD} \asymp K_1(\Phi\sqrt{\frac{\log p}{n+N}} + (\sigma + b_n + G_h)\sqrt{\frac{\log p}{n}})$, the estimator $\widehat{\theta}_{SD}$ in (3.5) achieves the following error bounds

$$\|\widehat{\theta}_{SD} - \theta^*\|_q = O_p\left(K_1 s^{1/q} \left\{ \Phi \sqrt{\frac{\log p}{n+N}} + (\sigma + b_n + G_h) \sqrt{\frac{\log p}{n}} \right\} \right), \quad (3.7)$$

for $q = 1, 2$. Moreover, if $G_h = 0$, that is, $f = h$, $b_n/\sigma = o(1)$ and $K_1 = O(1)$, we obtain

$$\|\widehat{\theta}_{SD} - \theta^*\|_q = O_p\left(s^{1/q} \left\{ \Phi \sqrt{\frac{\log p}{n+N}} + \sigma \sqrt{\frac{\log p}{n}} \right\} \right). \quad (3.8)$$

The error bound (3.7) depends on b_n , the stochastic error from the estimator \widehat{h}^{-j} , and G_h that can be viewed as the asymptotic bias of the estimator. If we assume f belongs to $\mathcal{F}_{\text{pairwise}}$ or $\mathcal{F}_{\text{additive}}$ introduced in Section 2, there exist estimators \widehat{h}^{-j} in the literature such that $G_h = 0$ and $b_n/\sigma = o(1)$; see further explanations in Remark 3.3 and more detailed examples in Supplement S.5. Thus, with $K_1 = O(1)$, the error bound (3.7) reduces to (3.8), which matches the minimax lower bound with $q = 1, 2$ in Theorem 2.2 up to a $\log s$ factor. Therefore, we call the estimator under this situation the optimal semi-supervised estimator. On the other hand, if the bias term G_h is large, the error bound (3.7) implies that our estimator $\widehat{\theta}_{SD}$ may have a slow rate. We will revisit this problem in Section 4.

Practically, we can set the tuning parameter $\lambda_{SD} = CK_1(\widehat{\Phi}\sqrt{\frac{\log p}{n+N}} + \widehat{\sigma}\sqrt{\frac{\log p}{n}})$ and choose the constant C by cross-validation. To account for the magnitude of Φ^2 and σ^2 in λ_{SD} , we suggest to estimate Φ^2 and σ^2 by $\widehat{\Phi}^2 = (\widehat{\Phi}_1^2 + \widehat{\Phi}_2^2)/2$ and $\widehat{\sigma}^2 = (\widehat{\sigma}_1^2 + \widehat{\sigma}_2^2)/2$, where $\widehat{\Phi}_j^2 = \frac{1}{n_j} \sum_{i \in D_j} (\widehat{h}^{-j}(X_i) - \widehat{\theta}_D^T X_i)^2$ and $\widehat{\sigma}_j^2 = \frac{1}{n_j} \sum_{i \in D_j} (Y_i - \widehat{h}^{-j}(X_i))^2$. The cross-validation method works well in our simulations.

Remark 3.3. (Examples of \widehat{h}). If the knowledge of $f(X)$ is available to some extent, we can leverage this information to construct estimators of $f(X)$. We introduced two model classes $\mathcal{F}_{\text{pairwise}}$ and $\mathcal{F}_{\text{additive}}$ in Section 2. For instance, if $f(\cdot) \in \mathcal{F}_{\text{additive}}$, we can directly apply the existing estimators in the literature; see Lin and Zhang (2006), Meier, Van de Geer, and Bühlmann (2009), Huang, Horowitz, and Wei (2010), and Raskutti, Wainwright, and Yu (2012), among many others. In particular, corollary 2 in Huang, Horowitz, and Wei (2010) implies that their adaptive group lasso estimator \widehat{h} constructed with spline basis functions satisfies $\|\widehat{h} - f\|_2 = O_p(n^{-d/(2d+1)})$, where d is the smoothness of the function $f_k(\cdot)$.

Remark 3.4. (Comparison with Alquier and Hebiri 2012). In the semi-supervised setting, Alquier and Hebiri (2012) proposed a transductive version of lasso and Dantzig selector based on the imputation method. Let $\tilde{Y} = (\tilde{Y}_1, \dots, \tilde{Y}_{n+N})$ denote the imputed outcomes (or pseudo-outcomes) from some preliminary estimator. The transductive Dantzig selector is given by

$$\hat{\theta}_T = \operatorname{argmin} \|\theta\|_1 \text{ s.t. } \frac{1}{n+N} \|\tilde{X}(\tilde{Y} - \tilde{X}^T \theta)\|_\infty \leq \lambda_T. \quad (3.9)$$

If the imputation satisfies $\frac{1}{n+N} \|\tilde{X}(\tilde{Y} - \tilde{X}\theta^*)\|_\infty \lesssim (\sigma + \Phi)\sqrt{\frac{\log p}{n}}$, it can be shown that $\|\hat{\theta}_T - \theta^*\|_1 \lesssim (\sigma + \Phi)s\sqrt{\frac{\log p}{n}}$ with high probability. We can see that the error bound is of the same order as the supervised Dantzig selector (2.3). Thus, the transductive Dantzig selector is also sub-optimal in the minimax sense; see the last paragraph of Section 2.

To see how the transductive Dantzig selector differs from our estimator $\hat{\theta}_{SD}$, we assume that the imputation is attained by using our estimator $\hat{h}(\cdot)$, that is, $\tilde{Y}_i = \hat{h}(X_i)$ for $1 \leq i \leq n+N$. For simplicity, let us omit the cross-fitting step here and defer more derivations in Supplement S.4.3. Then the modified score function $\hat{\Sigma}_{n+N}\theta - \hat{\xi}$ in our estimator (3.5) can be rewritten as

$$\begin{aligned} \hat{\Sigma}_{n+N}\theta - \hat{\xi} &= \frac{1}{n+N} \sum_{i=1}^n X_i(Y_i + \frac{N}{n}\{Y_i - \tilde{Y}_i\} - X_i^T \theta) \\ &\quad + \frac{1}{n+N} \sum_{i=n+1}^{n+N} X_i(\tilde{Y}_i - X_i^T \theta). \end{aligned}$$

We can see that for the samples from the unlabeled data ($i \in \{n+1, \dots, n+N\}$), we use \tilde{Y}_i as the pseudo-outcome. However, for the samples from the labeled data ($i \in \{1, \dots, n\}$), $Y_i + \frac{N}{n}\{Y_i - \tilde{Y}_i\}$ serves as the pseudo-outcome. In comparison, the transductive Dantzig selector (3.9) always uses \tilde{Y}_i as the pseudo-outcome for $i \in \{1, \dots, n+N\}$.

Remark 3.5. (Comparison with Bellec et al. 2018). Recently, Bellec et al. (2018) proposed a modified lasso estimator for prediction in the semi-supervised setting, which can be reformatted as the following Dantzig selector

$$\hat{\theta}_U = \arg \min \|\theta\|_1, \text{ s.t. } \|\hat{\Sigma}_{n+N}\theta - \frac{1}{n} \sum_{i=1}^n X_i Y_i\|_\infty \leq \lambda_U, \quad (3.10)$$

where $\hat{\Sigma}_{n+N} = \frac{1}{n+N} \sum_{i=1}^{n+N} X_i^{\otimes 2}$. Compared with our estimator $\hat{\theta}_{SD}$ in (3.5), $\hat{\theta}_U$ turns out to be a special case of $\hat{\theta}_{SD}$ by plugging $h = 0$ into (3.4).

Bellec et al. (2018) showed that if a large number of unlabeled data are used to compute $\hat{\Sigma}_{n+N}$, it becomes more plausible to assume that the compatibility (or RE) constant is bounded away from zero. The statement also holds true for our semi-supervised estimator $\hat{\theta}_{SD}$ with the use of $\hat{\Sigma}_{n+N}$. Moreover, they proved that the error bound for the excess risk in prediction remains $O_p(s \log p/n)$ under certain conditions, including $|Y| \leq C$ for some constant $C > 0$ which indeed implies $\Phi = O(1)$ and $\sigma = O(1)$ by their proof of Theorem 7. To make a fair comparison of $\hat{\theta}_U$ with our estimator $\hat{\theta}_{SD}$, we show that under the same conditions in our Theorem 3.2,

$$\|\hat{\theta}_U - \theta^*\|_1 = O_p\left(s(\Phi + \sigma + (\theta^{*T} \Sigma \theta^*)^{1/2})\sqrt{\frac{\log p}{n}}\right). \quad (3.11)$$

The proof is deferred to Supplement S.4.2. It is seen that $\hat{\theta}_U$ has a slower rate than our estimator $\hat{\theta}_{SD}$ if $\Phi/\sigma \rightarrow \infty$ or $\theta^{*T} \Sigma \theta^*/\sigma^2 \rightarrow \infty$. Again, in Supplement S.4.1, we consider the examples where $\Phi^2 \asymp s$ and σ is a constant, so that $\Phi/\sigma \rightarrow \infty$ holds, as s grows with n . Perhaps, a more surprising fact is that the convergence rate of $\hat{\theta}_U$ can be even slower than the fully supervised Dantzig selector $\hat{\theta}_D$ in (2.3) if $\theta^{*T} \Sigma \theta^*/(\sigma^2 + \Phi^2) \rightarrow \infty$. Indeed, our simulation studies confirm that the estimator $\hat{\theta}_U$ often produces larger estimation error than $\hat{\theta}_{SD}$ and $\hat{\theta}_D$.

4. Safe Semi-Supervised Estimator

Recall from Theorem 3.2 that our semi-supervised estimator $\hat{\theta}_{SD}$ is minimax optimal, if the conditional mean function $f(\cdot)$ can be consistently estimated with a proper rate, for example, when $f(\cdot)$ belongs to $\mathcal{F}_{\text{pairwise}}$ or $\mathcal{F}_{\text{additive}}$ introduced in Section 2. If this does not hold, there is no guarantee that the estimator $\hat{\theta}_{SD}$ attains the minimax lower bound or outperforms the supervised estimator. In particular, when $G_h + b_n \gg \Phi$, the convergence rate of $\hat{\theta}_{SD}$ can be even slower than the supervised estimator $\hat{\theta}_D$ in (2.2), hence, the integration of unlabeled data might fail to improve the estimation accuracy of θ^* .

To tackle this problem, we develop a safe semi-supervised estimator via a two-step procedure to adapt to the unknown quality of the conditional mean model $h(\cdot)$, which makes the final estimator no worse than the supervised estimator. To facilitate the theoretical analysis of the two-step estimator, we proceed with the lasso type estimators in this section. Recall that the supervised lasso was defined in (2.4) in Section 2. Equivalent to (3.6), we can rewrite our optimal semi-supervised lasso estimator $\hat{\theta}_{SL}$ as

$$\begin{aligned} \hat{\theta}_{SL} = \operatorname{argmin}_{\theta \in \mathbb{R}^p} & \sum_{j=1}^2 \frac{\sum_{i \in D_j} (\hat{h}^{-j}(X_i) - X_i^T \theta)^2}{n+N} \\ & - \frac{2 \sum_{i \in D_j^*} (Y_i - \hat{h}^{-j}(X_i)) X_i^T \theta}{n} + 2\lambda_{SL} \|\theta\|_1. \end{aligned} \quad (4.1)$$

We construct the safe semi-supervised estimator in the following two steps. First, for a given estimate of the unknown conditional mean \hat{h} , we compute our semi-supervised lasso estimator $\hat{\theta}_{SL}$ in (4.1) and the supervised lasso estimator $\hat{\theta}_L$ in (2.4). Since the estimator $\hat{\theta}_{SL}$ may not be desirable when the corresponding conditional mean model h is misspecified, in the second step we further improve the initial estimator $\hat{\theta}_{SL}$ by a step of refitting on a suitable estimate of the support set of θ^* . Denote the support set of two lasso type estimators $\hat{\theta}_L$ and $\hat{\theta}_{SL}$ by \hat{T}_1 and \hat{T}_2 , respectively. We define the safe semi-supervised estimator as

$$\hat{\theta}_R = \hat{\theta}_{SL} + \hat{\omega}, \quad (4.2)$$

where $\hat{\omega} \in \mathbb{R}^p$ is attained by

$$\hat{\omega} = \operatorname{argmin}_{\operatorname{supp}(\omega) \subseteq \hat{T}_1 \cup \hat{T}_2} \frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T (\hat{\theta}_{SL} + \omega))^2 + \lambda_\omega \|\omega\|_1. \quad (4.3)$$

In (4.3) we impose $\operatorname{supp}(\omega) \subseteq \hat{T}_1 \cup \hat{T}_2$, which implies $\hat{\omega}_j = 0$ for any $j \notin \hat{T}_1 \cup \hat{T}_2$. Thus, the safe semi-supervised estimator $\hat{\theta}_R$ may only differ from $\hat{\theta}_{SL}$ on the estimated support set $\hat{T}_1 \cup \hat{T}_2$.

Since both the estimands of $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_R$ are $\boldsymbol{\theta}^*$, by the definition of (4.2) we can treat $\widehat{\boldsymbol{\omega}}$ as an estimator of $\boldsymbol{\theta}^* - \boldsymbol{\theta}^* = 0$. At a first sight, one may expect that estimating a known vector of 0 by $\widehat{\boldsymbol{\omega}}$ and summing with $\widehat{\boldsymbol{\theta}}_{SL}$ will inflate the error of the resulting estimator. However, we show a surprising result that the two step estimator $\widehat{\boldsymbol{\theta}}_R$ may outperform the initial estimator $\widehat{\boldsymbol{\theta}}_{SL}$.

In the following, we will first explain the intuition behind the estimator $\widehat{\boldsymbol{\theta}}_R$. Let us consider two scenarios. First, if our initial estimator $\widehat{\boldsymbol{\theta}}_{SL}$ in step 1 is able to achieve a fast rate, one would expect that $\widehat{\boldsymbol{\omega}} \approx 0$ due to the L_1 regularization in (4.3). Thus, $\widehat{\boldsymbol{\theta}}_R$ will inherit the fast rate from the initial estimator $\widehat{\boldsymbol{\theta}}_{SL}$. In the second scenario, if $\widehat{\boldsymbol{\theta}}_{SL}$ has a slow rate due to model misspecification, with a suitable choice of tuning parameters, we can obtain $\widehat{\boldsymbol{\omega}} \approx \widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}$. As a toy example, if p is fixed and small and we set all tuning parameters in $\widehat{\boldsymbol{\omega}}$, $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$ to be 0, it is easily shown from the least square formula that $\widehat{\boldsymbol{\omega}} = \widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}$. Thus, in this case, the safe semi-supervised estimator will resemble the supervised estimator, since $\widehat{\boldsymbol{\theta}}_R = \widehat{\boldsymbol{\theta}}_{SL} + \widehat{\boldsymbol{\omega}} \approx \widehat{\boldsymbol{\theta}}_L$. In summary, the refitting step can simultaneously retain the fast rate of $\widehat{\boldsymbol{\theta}}_{SL}$ if it has, and alleviate the potentially unsatisfactory performance of $\widehat{\boldsymbol{\theta}}_{SL}$ under model misspecification. The following theorem shows the convergence rate of the safe semi-supervised estimator $\widehat{\boldsymbol{\theta}}_R$.

Theorem 4.1. Suppose Assumption 3.1 holds, $\Lambda_{\max}(\boldsymbol{\Sigma}) \leq C_{\max} < \infty$ and $s(\log p)^2 = o(n)$. Assume that we choose the tuning parameters λ_{SL} , λ_L and λ_ω in (4.1), (2.4), and (4.3) as $\lambda_{SL} \asymp K_1(\Phi\sqrt{\frac{\log p}{n+N}} + (\sigma + b_n + G_h)\sqrt{\frac{\log p}{n}})$ and $\lambda_L \asymp \lambda_\omega \asymp K_1(\Phi + \sigma)\sqrt{\frac{\log p}{n}}$. We have for $q = 1, 2$,

$$\|\widehat{\boldsymbol{\theta}}_R - \boldsymbol{\theta}^*\|_q = O_p(R_{SL} \wedge R_L),$$

where $R_{SL} = s^{1/q} K_1 \left(\Phi \sqrt{\frac{\log p}{n+N}} + (\sigma + b_n + \|f - h\|_2) \sqrt{\frac{\log p}{n}} \right)$, and $R_L = s^{1/q} K_1(\Phi + \sigma)\sqrt{\frac{\log p}{n}}$.

Recall that R_{SL} and R_L correspond to the convergence rates of $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$, respectively. Theorem 4.1 shows that $\widehat{\boldsymbol{\theta}}_R$ attains the faster rate between $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$. Thus, the estimator $\widehat{\boldsymbol{\theta}}_R$ remains minimax rate-optimal when the initial estimator $\widehat{\boldsymbol{\theta}}_{SL}$ is optimal (see Theorem 3.2), and is guaranteed to be no worse than the supervised estimators even if the conditional mean model is misspecified.

To choose the tuning parameters, we note that the magnitude of λ_{SL} is the same as λ_{SD} in Theorem 3.2. We can apply the same cross-validation method explained after Theorem 3.2 to tune λ_{SL} . Since λ_L can be written as $\lambda_L = CK_1\{\mathbb{E}(Y - X^T\boldsymbol{\theta}^*)^2\}^{1/2}\sqrt{\frac{\log p}{n}}$ for some constant C , one may use scaled lasso to estimate the noise level $\mathbb{E}(Y - X^T\boldsymbol{\theta}^*)^2$, and further apply cross-validation to tune λ_L . In practice, to reduce the computational cost, once we have selected the tuning parameter λ_L , we can simply set $\lambda_\omega = \lambda_L$, which works well in our simulations.

Remark 4.2. The two-step approach is inspired by the recent works of Bastani (2020) and Li, Cai, and Li (2020) in the context of transfer learning. However, our theoretical guarantee in Theorem 4.1 is much stronger than those works. Using our terminology, the theoretical analysis in Bastani (2020) and Li, Cai, and Li (2020) implied that the estimator $\widehat{\boldsymbol{\theta}}_R$ can only attain

the worst possible rate between $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$, that is, $\|\widehat{\boldsymbol{\theta}}_R - \boldsymbol{\theta}^*\|_q = O_p(R_{SL} + R_L)$, which is not sufficient to show $\widehat{\boldsymbol{\theta}}_R$ to be safe. We obtain a sharper result in Theorem 4.1 because our refitting step (4.3) differs from those works. In particular, unlike their works, we constrain the support of the estimator $\widehat{\boldsymbol{\omega}}$ to be $\widehat{T}_1 \cup \widehat{T}_2$, which guarantees the sparsity of $\widehat{\boldsymbol{\theta}}_R$. This is an important intermediate step to prove Theorem 4.1. Finally, we note that, in the context of transfer learning, Li, Cai, and Li (2020) also considered how to use model aggregation, such as Q-aggregation (Dai, Rigollet, and Zhang 2012), to improve the initial estimator. We refer to Supplement S.4.4 for further discussion on model aggregation.

5. Aggregation of Semi-Supervised Estimators

Recall from Theorem 3.2 that the performance of the semi-supervised estimator depends on the estimator \widehat{h} of the conditional mean model. In practice, it is uncommon for us to find a consistent estimator of the conditional mean function. More commonly, we may face multiple choices of \widehat{h} 's, none of which is consistent. We may expect that different \widehat{h} 's only capture different aspects of the non-linearity of the conditional mean function. Does aggregating multiple semi-supervised estimators help us explore the non-linearity of the conditional mean function? Here, we generalize the two-step method in Section 4 to answer this question.

Assume that two different estimators of the conditional mean function \widehat{h}_1 and \widehat{h}_2 are available. We denote by $\widehat{\boldsymbol{\theta}}_{h_1}$ and $\widehat{\boldsymbol{\theta}}_{h_2}$ the semi-supervised lasso estimators in (4.1) with \widehat{h}_1 and \widehat{h}_2 and tuning parameters λ_{h_1} and λ_{h_2} , respectively. In the following, we will apply the two-step procedure to combine $\widehat{\boldsymbol{\theta}}_{h_1}$ and $\widehat{\boldsymbol{\theta}}_{h_2}$. We first compute $\widehat{\boldsymbol{\theta}}_{h_1}$ and $\widehat{\boldsymbol{\theta}}_{h_2}$ in step 1. Their support sets are denoted by \widehat{H}_1 and \widehat{H}_2 . In step 2, we use $\widehat{\boldsymbol{\theta}}_{h_1}$ as the initial estimator and define the aggregated estimator as

$$\widehat{\boldsymbol{\theta}}_{AH} = \widehat{\boldsymbol{\theta}}_{h_1} + \widehat{\boldsymbol{\omega}}_h, \quad (5.1)$$

where

$$\widehat{\boldsymbol{\omega}}_h = \underset{\text{supp}(\boldsymbol{\omega}) \subseteq \widehat{H}_1 \cup \widehat{H}_2}{\text{argmin}} \sum_{j=1}^2 \left(\frac{\sum_{i \in D_j} (\widehat{h}_2^{-j}(X_i) - X_i^T(\widehat{\boldsymbol{\theta}}_{h_1} + \boldsymbol{\omega}))^2}{n+N} \right. \\ \left. - \frac{2 \sum_{i \in D_j^*} (Y_i - \widehat{h}_2^{-j}(X_i)) X_i^T(\widehat{\boldsymbol{\theta}}_{h_1} + \boldsymbol{\omega})}{n} \right) \\ + 2\lambda_H \|\boldsymbol{\omega}\|_1.$$

Unlike the refitting step in (4.3), we also incorporate the unlabeled data to compute $\widehat{\boldsymbol{\omega}}_h$ so that $\widehat{\boldsymbol{\theta}}_{AH}$ can attain a better convergence rate than $\widehat{\boldsymbol{\theta}}_{h_1}$ or $\widehat{\boldsymbol{\theta}}_{h_2}$. One may note that the creation of $\widehat{\boldsymbol{\theta}}_{AH}$ is not symmetric to h_1 and h_2 ; however, our theoretical analysis below shows that the priority order of h_1 and h_2 does not really matter.

Proposition 5.1. Suppose Assumption 3.1 holds, $\Lambda_{\max}(\boldsymbol{\Sigma}) \leq C_{\max} < \infty$ and $s(\log p)^2 = o(n+N)$. The estimators \widehat{h}_1 and \widehat{h}_2 satisfy $\|\widehat{h}_1^{-j} - h_1\|_2 = O_p(b_{h_1})$ and $\|\widehat{h}_2^{-j} - h_2\|_2 = O_p(b_{h_2})$ for $j = 1, 2$. Denote $G_{h_1} = \|h_1 - f\|_2$ and $G_{h_2} = \|h_2 - f\|_2$. Selecting some tuning parameters $\lambda_{h_1} \asymp K_1(\Phi\sqrt{\frac{\log p}{n+N}} + (\sigma + b_{h_1} +$

$G_{h_1})\sqrt{\frac{\log p}{n}}, \lambda_{h_2} \asymp \lambda_H \asymp K_1(\Phi\sqrt{\frac{\log p}{n+N}} + (\sigma + b_{h_2} + G_{h_2})\sqrt{\frac{\log p}{n}})$, we can show that for $q = 1, 2$,

$$\|\widehat{\boldsymbol{\theta}}_{AH} - \boldsymbol{\theta}^*\|_q = O_p(R_{h_1} \wedge R_{h_2}),$$

where $R_{h_j} = s^{1/q}K_1\left(\Phi\sqrt{\frac{\log p}{n+N}} + (\sigma + b_{h_j} + \|f - h_j\|_2)\sqrt{\frac{\log p}{n}}\right)$.

This proposition shows that the aggregated estimator $\widehat{\boldsymbol{\theta}}_{AH}$ attains the best possible rate between $\widehat{\boldsymbol{\theta}}_{h_1}$ and $\widehat{\boldsymbol{\theta}}_{h_2}$. Following the same reasoning, one may use $\widehat{\boldsymbol{\theta}}_{AH}$ as the initial estimator and repeat the refitting step, if a third estimator $\widehat{\boldsymbol{\theta}}_{h_3}$ is available. The resulting estimator attains the best possible rate among the three estimators $\widehat{\boldsymbol{\theta}}_{h_1}$, $\widehat{\boldsymbol{\theta}}_{h_2}$, and $\widehat{\boldsymbol{\theta}}_{h_3}$. We expect that, in general, our two-step procedure can be applied to aggregate multiple estimators, as long as the number of the candidates is fixed and small. We refer to Section 6 for numerical performance of the aggregated estimators.

6. Simulation Studies

6.1. Data Generating Models and Practical Implementation

We first generate a p -dimensional multivariate normal random vector $Z \sim \mathcal{N}(0, \Sigma)$ with $\Sigma_{jk} = 0.3^{|j-k|}$. We set the covariate $X = (X_1, \dots, X_p)$ to be $X_1 = |Z_1|$ and $X_j = Z_j$ for $1 < j \leq p$. The reason we take $X_1 = |Z_1|$ is that this transformation implies $\mathbb{E}(X_1^k X_j) = 0$ for $j \neq 1$ but the parameter θ_1^* for centered X_1 is nonzero. We consider the following three data generating models for Y . For Model 1, we consider an additive model

$$Y = 0.5X_1^2 + 0.8X_3^3 - (X_4 - 2)^2 + 2(X_5 + 1)^2 + 2X_6 + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, 1)$. To calculate the corresponding regression parameter $\boldsymbol{\theta}^*$ under the working linear model, we first center Y and X_1 so that their means are 0. By Proposition 4 in Bühlmann and Van de Geer (2015), we know that the support of $\boldsymbol{\theta}^*$ is $S = \{1, 3, 4, 5, 6\}$ and θ_j^* for any $j \in S$ is given by the $L_2(\mathbb{P})$ projection in the sub-model only with the variable X_j (e.g., $\theta_3^* = \arg \min \mathbb{E}(0.8X_3^2 - \theta_3 X_3)^2$). After some calculation, we obtain $\boldsymbol{\theta}^* = (1.1, 0, 2.4, 4, 4, 2, 0, \dots, 0)$, which is sparse.

For Model 2, we consider

$$Y = 0.6(X_1 + X_2)^2 + 0.4X_4^3 - X_5 + 2X_6 + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, 1)$. The model is non-additive since it includes an interaction term between X_1 and X_2 . The corresponding regression parameter $\boldsymbol{\theta}^*$ is $(1.48, 1.04, 0, 1.2, -1, 2, 0, \dots, 0)$.

Besides, we consider a nonadditive Model 3 which includes extra randomly selected support sets S_1 and S_2 on top of Model 1. We choose $|S_1| = |S_2| = 5$ and set

$$Y = 0.5X_1^2 + 0.8X_3^3 - (X_4 - 2)^2 + 2(X_5 + 1)^2 + 2X_6 + 2 \sum_{k \in S_1} X_k + 0.05(\sum_{k \in S_2} X_k)^3 + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, 1)$. In this model, $\theta_k^* = 0.75$ for $k \in S_2$, $\theta_k^* = 2$ for $k \in S_1$, and the other components of $\boldsymbol{\theta}^*$ are the same as in Model 1. Under each data generating model, we consider several combinations of (n, p) and vary the ratio N/n in a certain range.

We repeat the simulation 100 times and output the average as final results.

Before we proceed to illustrate the results, we list the estimators considered in this section and discuss several practical implementation issues. Since the performance of lasso and Dantzig type estimators are similar, we only consider the Dantzig type estimators here for simplicity.

- The proposed semi-supervised estimator $\widehat{\boldsymbol{\theta}}_{SD}$ in (3.5) (SSL1) with a sparse additive model h_1 . We estimate the sparse additive model by using the group lasso with the spline basis (Huang, Horowitz, and Wei 2010). To be specific, we use the cubic spline basis with degree of freedom $df = 5$. To select the penalty parameter in group lasso and make computation easier, the BIC criterion is used; see Section 4 in Huang, Horowitz, and Wei (2010) for the definition. After we derive the estimator \widehat{h}_1 and subsequently $\widehat{\xi}$, we modify the source code in the `flare` package to compute $\widehat{\boldsymbol{\theta}}_{SD}$, where the tuning parameter λ_{SD} is selected by 5-fold cross-validation.
- The proposed semi-supervised estimator $\widehat{\boldsymbol{\theta}}_{SD}$ in (3.5) (SSL2) with a pairwise interaction model h_2 . The model $h_2(\cdot)$ corresponds to the linear regression containing all the linear terms, the squares of the variable and the interaction terms (Zhao and Leng 2016).
- The supervised Dantzig selector $\widehat{\boldsymbol{\theta}}_D$ in (2.2) (Dantzig). We use the `flare` package to compute the estimator and select the tuning parameter by 5-fold cross-validation.
- The modified Dantzig selector $\widehat{\boldsymbol{\theta}}_U$ in (3.10) (U-Dantzig).
- The aggregated estimator $\widehat{\boldsymbol{\theta}}_{AH}$ in (5.1) (SSL12) that combines SSL1 and SSL2.
- The safe semi-supervised estimator $\widehat{\boldsymbol{\theta}}_R$ (S-SSL1) using SSL1 as the initial estimator. As seen from the discussion after Theorem 4.1, we set the tuning parameter λ_ω identical to the tuning parameter for the supervised Dantzig selector.
- The safe semi-supervised estimator $\widehat{\boldsymbol{\theta}}_R$ (S-SSL12) using SSL12 (e.g., the aggregated estimator) as the initial estimator.

6.2. Numerical Results

Under Model 1 with $p = 500$ and $n = 200$, the comparisons of the estimation errors of SSL1, SSL2, Dantzig, SSL12, S-SSL1, and S-SSL12 are illustrated in Figure 1, while the results from U-Dantzig are shown in numbers in the caption of the figure due to the scale of plot. Since the true data generating model is additive with some quadratic terms, both SSL1 and SSL2 can leverage the non-linearity of the conditional mean function and their estimation errors are much smaller than Dantzig, which agrees with Theorem 3.2. Besides, we can see that SSL1 outperforms SSL2, since the imposed additive model in SSL1 can better estimate the conditional mean function. By aggregating SSL1 and SSL2, our estimator SSL12 achieves the minimum L_2 error among all those methods. For the two safe semi-supervised estimators (S-SSL1 and S-SSL12), they retain the optimal rate in L_2 norm from the corresponding semi-supervised estimators (SSL1 and SSL12), and clearly outperform Dantzig, which is consistent with the theoretical property in Theorem 4.1.

One interesting observation is that U-Dantzig performs much worse than the fully supervised estimator Dantzig; see the caption of Figure 1. Thus, using the sample covariance $\widehat{\Sigma}_{n+N}$

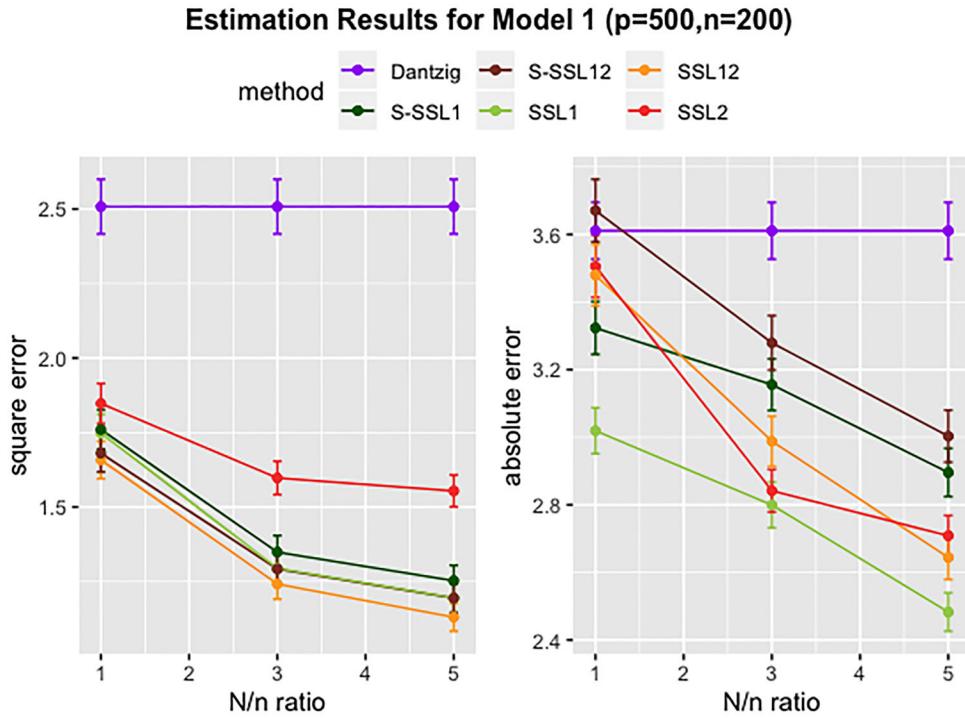


Figure 1. The L_2 and L_1 estimation errors under Model 1 with $p = 500$ and $n = 200$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 5.75 (0.20), 7.14 (0.24), 7.47 (0.26) and L_1 errors for U-Dantzig are 6.15 (0.17), 6.53 (0.15), 6.73 (0.15). The numbers in the parenthesis are sample standard deviations.

from both labeled and unlabeled data in the Dantzig selector may not provide any empirical improvement; see Remark 3.5 for the theoretical justification. In addition, as the size of unlabeled data N increases, the improvement of our semi-supervised estimators (SSL1, SSL2, SSL12, S-SSL1, and S-SSL12) is more overwhelming, whereas the performance of U-Dantzig tends to deteriorate.

The comparisons under Model 2 with $p = 500$ and $n = 200$ are shown in Figure 2. Since Model 2 includes an interaction term between X_1 and X_2 , the sparse additive model h_1 is inconsistent for the true regression function. Thus, SSL1 does not improve the estimation accuracy compared with fully supervised Dantzig. However, the safe semi-supervised estimator, S-SSL1, successfully mitigates the undesired performance of SSL1 and its L_1 and L_2 errors are smaller than Dantzig. This agrees with Theorem 4.1 that the refitting step provides a safe use of unlabeled data even if the imposed conditional mean model is incorrect. On the other hand, since the pairwise interaction model h_2 is a correctly specified conditional mean model, the estimators SSL2, SSL12, and S-SSL12, that depend on this model, show small estimation errors.

Under Model 3, since the true conditional mean function differs significantly from the additive model h_1 , SSL1 yields large estimation errors. Nevertheless, the performance of S-SSL1 is comparable and no worse than the fully supervised Dantzig. While the pairwise interaction model h_2 cannot account for the third order interaction terms in the set S_2 , it can still partially explain the non-linearity of the true conditional mean function. Thus, the performance of SSL2 is still better than Dantzig in L_2 norm when $N/n = 5$ and also in L_1 norm. The comparison results are summarized in Figure 3.

The Supplement S.6 contains further simulation results with $p = 200$ and $n = 100$, and with $p = 1000$ and $n = 300$, for all the three models considered above, as well as some other numerical results.

7. Real Data Application

In this section, we illustrate our proposed methods in a real data example, derived from the Medical Information Mart for Intensive Care III (MIMIC-III) database (Johnson et al. 2016). MIMIC-III is an openly available electronic health records system developed by the MIT Lab for Computational Physiology. It contains de-identified health-related data for 38,597 adult patients (aged 16 years or above) admitted to intensive care units (ICU) of the Beth Israel Deaconess Medical Center between 2001 and 2012. Some details of the adult patients by their first ICU admissions are available in the Table 1 of Johnson et al. (2016).

Researchers have used the MIMIC-III database to investigate a variety of medical issues such as predicting ICU readmissions (Brown et al. 2012; Tabak et al. 2017; Xue, Klabjan, and Luo 2019) and associating various clinical biomarkers with mortality (Liu et al. 2020; Du et al. 2021; Jhou et al. 2021; Tang et al. 2021). Our initial motivation for this data analysis is to understand the relation between the albumin level in the blood sample, oftentimes used to screen for liver or kidney disease (Phillips, Shaper, and Whincup 1989), and all other variables including demographics, chart events, and clinical biomarkers from the laboratory tests.

After all of the data pre-processing steps detailed in Supplement S.7, we are left with a dataset with 4784 patients and each of

Estimation Results for Model 2 (p=500,n=200)

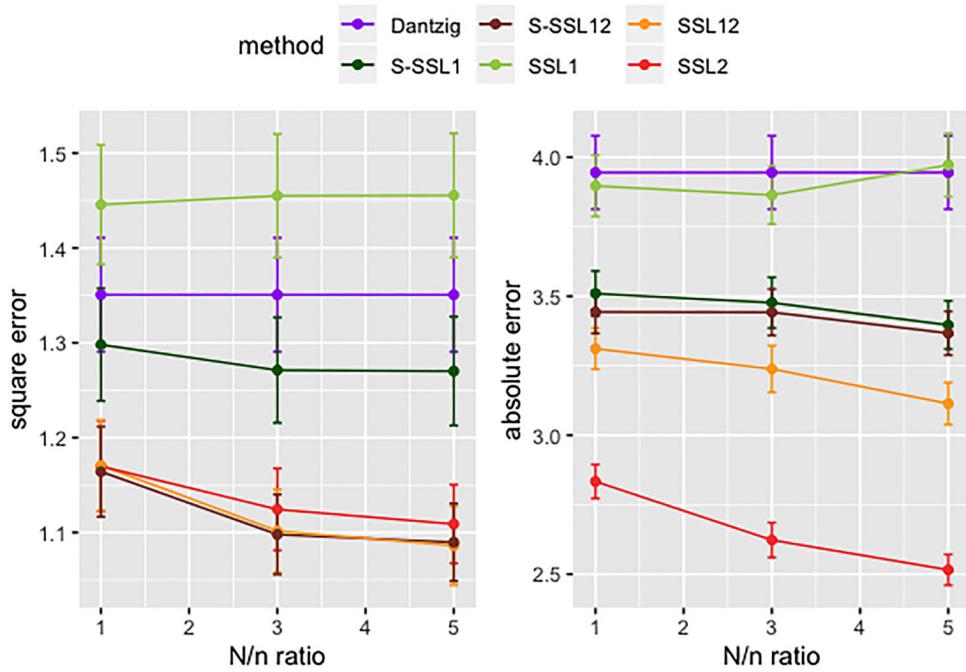


Figure 2. The L_2 and L_1 estimation errors under Model 2 with $p = 500$ and $n = 200$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 2.31 (0.08), 2.75 (0.09), 2.87 (0.11) and L_1 errors for U-Dantzig are 4.28 (0.10), 4.62 (0.10), 4.66 (0.07).

Estimation Results for Model 3 (p=500,n=200)

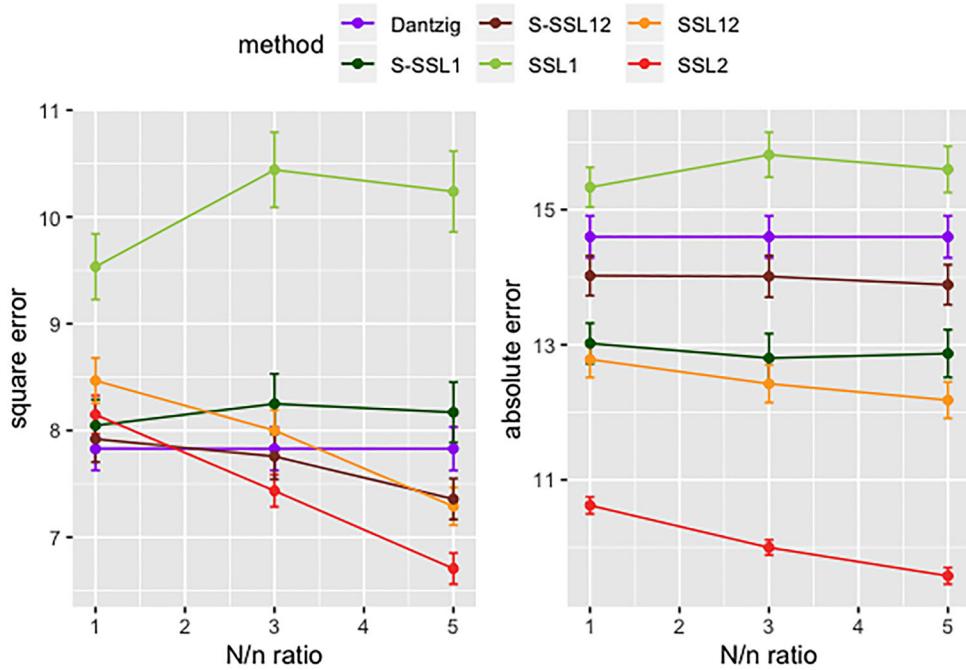


Figure 3. The L_2 and L_1 estimation error under Model 3 with $p = 500$ and $n = 200$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 11.46(0.23), 13.29(0.25), 13.62(0.26) and L_1 errors for U-Dantzig are 14.32(0.16), 15.73(0.24), 15.60(0.20).

them has $p = 2928$ covariates. The results depicted in Figure 4 are an aggregation from 100 replications. In each replication, we randomly choose 2000 patients and call the first $n = 500$ patients the labeled data. We mask the outcome “albumin” for all other 1500 patients and call them the unlabeled data. In each experiment, we gradually increase the sample size of the unlabeled data from $N = 500$ to $N = 1000$ and to $N = 1500$. Since

the working model already includes some pairwise interaction terms among biomarkers, we use the random forest to estimate the conditional mean function in our SSL estimator and the corresponding S-SSL estimator. The L_1 and L_2 estimation errors presented in Figure 4 are computed against the lasso estimator using all of the 4784 samples, which is regarded as the proxy of the underlying true linear coefficients of the working model.

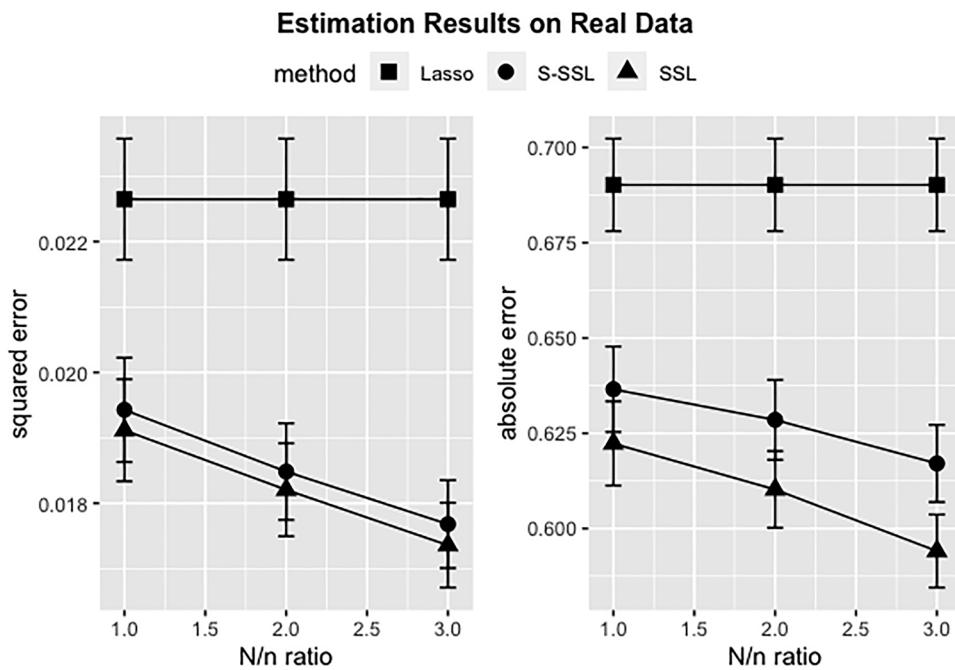


Figure 4. The L_2 and L_1 estimation error for real data application. The length of the vertical bar represents the magnitude of the sample standard deviations.

From Figure 4, as the ratio N/n increases from 1 to 3, compared to supervised lasso, the proposed S-SSL estimator could bring the L_1 error down around 7.6% to 10.3% and the proposed SSL estimator could bring the L_1 error down around 10.1% to 21.7%, respectively. The percentage decrease for the L_2 error ranges from 14.9% to 22.4% for the S-SSL estimator and from 23.2% to 28.6% for the SSL estimator. Compared to the S-SSL estimator, the outperformance of the SSL estimator, albeit not as significant as the comparison with the supervised lasso, is still noticeable from Figure 4. The likely reason for this is that the random forest used in our semi-supervised estimators (SSL and S-SSL) has already effectively captured the structure of the conditional mean function. Thus, S-SSL behaves similarly to SSL. All of these observations support the effectiveness of the methods proposed in this article. In Supplement S.7, similar results can be obtained where the estimation errors are computed against the lasso which exclusively employs samples that are not used in the computation of the three comparators.

Finally, Supplement S.7 also contains the detailed data cleaning and data pre-processing procedures for this application, as well as some other results and conclusions. Along the article, we also submit the programming code for anyone who has interest to reproduce the results.

Supplementary Materials

The supplementary materials include the technical proofs, some more detailed theoretical results and discussions, and additional numerical results.

Acknowledgments

The authors would like to thank the Editor, an Associate Editor, and three reviewers for their insightful comments which have helped improve the manuscript substantially.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Ning is supported in part by U.S. National Science Foundation (NSF DMS 1854637 and 1941945) and U.S. National Institutes of Health (NIH 1RF1AG077820-01A1). Zhao is supported in part by NSF (DMS 1953526, 2122074, and 2310942), NIH (R01DC021431) and the American Family Funding Initiative of UW-Madison. Zhang is supported in part by NIH (R01HG010171 and R01MH116527) and NSF (DMS 2112711).

References

- Abdullah, S. S., Rostamzadeh, N., Sedig, K., Garg, A. X., and McArthur, E. (2020), “Visual Analytics for Dimension Reduction and Cluster Analysis of High Dimensional Electronic Health Records,” in *Informatics* (Vol. 7), p. 17, Basel: Multidisciplinary Digital Publishing Institute. [1]
- Alquier, P., and Hebiri, M. (2012), “Transductive Versions of the LASSO and the Dantzig Selector,” *Journal of Statistical Planning and Inference*, 142, 2485–2500. [2,6]
- Azriel, D., Brown, L. D., Sklar, M., Berk, R., Buja, A., and Zhao, L. (2016), “Semi-Supervised Linear Regression,” arXiv e-prints, arXiv:1612.02391. [2]
- Bastani, H. (2020), “Predicting with Proxies: Transfer Learning in High Dimension,” *Management Science*, 67, 2657–3320. [7]
- Bellec, P. C., Dalalyan, A. S., Grappin, E., and Paris, Q. (2018), “On the Prediction Loss of the Lasso in the Partially Labeled Setting,” *Electronic Journal of Statistics*, 12, 3443–3472. [2,6]
- Bellec, P. C., Lecué, G., and Tsybakov, A. B. (2018), “Slope Meets Lasso: Improved Oracle Bounds and Optimality,” *The Annals of Statistics*, 46, 3603–3642. [3]
- Berk, R., Buja, A., Brown, L., George, E., Kuchibhotla, A., Su, W., and Zhao, L. (2019), “Assumption Lean Regression,” *The American Statistician*, 75, 76–84. [1]
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., and Rafel, C. A. (2019), “Mixmatch: A Holistic Approach to Semi-Supervised Learning,” in *Advances in Neural Information Processing Systems* (Vol. 32). [2]

Bickel, P. J. (1982), "On Adaptive Estimation," *The Annals of Statistics*, 10, 647–671. [4]

Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009), "Simultaneous Analysis of Lasso and Dantzig Selector," *The Annals of Statistics*, 37, 1705–1732. [4,5]

Brown, S. E., Ratcliffe, S. J., Kahn, J. M., and Halpern, S. D. (2012), "The Epidemiology of Intensive Care Unit Readmissions in the United States," *American Journal of Respiratory and Critical Care Medicine*, 185, 955–964. [9]

Buja, A., Brown, L., Berk, R., George, E., Pitkin, E., Traskin, M., Zhang, K., and Zhao, L. (2019), "Models as Approximations I: Consequences Illustrated with Linear Regression," *Statistical Science*, 34, 523–544. [1]

Bühlmann, P., and Van de Geer, S. (2015), "High-Dimensional Inference in Misspecified Linear Models," *Electronic Journal of Statistics*, 9, 1449–1473. [5,8]

Candes, E., and Tao, T. (2007), "The Dantzig Selector: Statistical Estimation When p is Much Larger than n ," *The Annals of Statistics*, 35, 2313–2351. [1]

Chakrabortty, A., and Cai, T. (2018), "Efficient and Adaptive Linear Regression in Semi-Supervised Settings," *The Annals of Statistics*, 46, 1541–1572. [2]

Chapelle, O., Scholkopf, B., and Zien, A. (2006), *Semi-Supervised Learning* (Vol. 2), Cambridge, MA: MIT Press. [1]

— (2009), "Semi-Supervised Learning," *IEEE Transactions on Neural Networks*, 20, 542–542. [2]

Chen, Y., Wei, C., Kumar, A., and Ma, T. (2020), "Self-Training Avoids Using Spurious Features Under Domain Shift," in *Advances in Neural Information Processing Systems* (Vol. 33), pp. 21061–21071. [2]

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018), "Double/Debiased Machine Learning for Treatment and Structural Parameters," *The Econometrics Journal*, 21, C1–C68. [4]

Dai, D., Rigollet, P., and Zhang, T. (2012), "Deviation Optimal Learning Using Greedy Q-Aggregation," *The Annals of Statistics*, 40, 1878–1905. [7]

Du, H., Siah, K. T. H., Ru-Yan, V. Z., Teh, R., Tan, C. Y. E., Yeung, W., Scaduto, C., Bolongaita, S., Cruz, M. T. K., Liu, M., et al. (2021), "Prediction of In-Hospital Mortality of Clostridioides difficile Infection Using Critical Care Database: A Big Data-Driven, Machine Learning Approach," *BMJ Open Gastroenterology*, 8, e000761. [9]

Gensheimer, M. F., Henry, A. S., Wood, D. J., Hastie, T. J., Aggarwal, S., Dudley, S. A., Pradhan, P., Banerjee, I., Cho, E., and Ramchandran, K. (2019), "Automated Survival Prediction in Metastatic Cancer Patients Using High-Dimensional Electronic Medical Record Data," *JNCI: Journal of the National Cancer Institute*, 111, 568–574. [1]

Huang, J., Horowitz, J. L., and Wei, F. (2010), "Variable Selection in Nonparametric Additive Models," *The Annals of Statistics*, 38, 2282–2313. [3,5,8]

Jhou, H.-J., Chen, P.-H., Yang, L.-Y., Chang, S.-H., and Lee, C.-H. (2021), "Plasma Anion Gap and Risk of In-Hospital Mortality in Patients with Acute Ischemic Stroke: Analysis from the MIMIC-IV Database," *Journal of Personalized Medicine*, 11, 1004. [9]

Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. (2016), "MIMIC-III, A Freely Accessible Critical Care Database," *Scientific Data*, 3, 160035. [9]

Kohane, I. S. (2011), "Using Electronic Health Records to Drive Discovery in Disease Genomics," *Nature Reviews Genetics*, 12, 417–428. [1]

Li, S., Cai, T. T., and Li, H. (2020), "Transfer Learning for High-dimensional Linear Regression: Prediction, Estimation, and Minimax Optimality," *Journal of the Royal Statistical Society, Series B*, 84, 149–173. [7]

Liao, K. P., Cai, T., Gainer, V., Goryachey, S., Zeng-treitler, Q., Raychaudhuri, S., Szolovits, P., Churchill, S., Murphy, S., and Kohane, I. (2010), "Electronic Medical Records for Discovery Research in Rheumatoid Arthritis," *Arthritis Care & Research*, 62, 1120–1127. [1]

Lin, Y., and Zhang, H. H. (2006), "Component Selection and Smoothing in Multivariate Nonparametric Regression," *The Annals of Statistics*, 34, 2272–2297. [3,5]

Liu, X., Feng, Y., Zhu, X., Shi, Y., Lin, M., Song, X., Tu, J., and Yuan, E. (2020), "Serum Anion Gap at Admission Predicts All-Cause Mortality in Critically Ill Patients with Cerebral Infarction: Evidence From the MIMIC-III Database," *Biomarkers*, 25, 725–732. [9]

Meier, L., Van de Geer, S., and Bühlmann, P. (2009), "High-Dimensional Additive Modeling," *The Annals of Statistics*, 37, 3779–3821. [3,5]

Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., and Goodfellow, I. (2018), "Realistic Evaluation of Deep Semi-Supervised Learning Algorithms," in *Advances in Neural Information Processing Systems* (Vol. 31). [2]

Phillips, A., Shaper, A. G., and Whincup, P. (1989), "Association between Serum Albumin and Mortality from Cardiovascular Disease, Cancer, and Other Causes," *The Lancet*, 334, 1434–1436. [9]

Raskutti, G., Wainwright, M. J., and Yu, B. (2012), "Minimax-Optimal Rates For Sparse Additive Models Over Kernel Classes Via Convex Programming," *Journal of Machine Learning Research*, 13, 389–427. [3,5]

Rigollet, P. (2006), "Generalization Error Bounds in Semi-Supervised Classification under the Cluster Assumption," *Journal of Machine Learning Research*, 8, 1369–1392. [2]

Robins, J. M., Li, L., Mukherjee, R., Tchetgen, E. T., and van der Vaart, A. (2017), "Minimax Estimation of a Functional on a Structured High-Dimensional Model," *The Annals of Statistics*, 45, 1951–1987. [4]

Schick, A. (1986), "On Asymptotically Efficient Estimation in Semiparametric Models," *The Annals of Statistics*, 14, 1139–1151. [4]

Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C. A., Cubuk, E. D., Kurakin, A., and Li, C.-L. (2020), "Fixmatch: Simplifying Semi-Supervised Learning with Consistency and Confidence," *Advances in Neural Information Processing Systems*, 33, 596–608. [2]

Tabak, Y. P., Sun, X., Nunez, C. M., Gupta, V., and Johannes, R. S. (2017), "Predicting Readmission at Early Hospitalization Using Electronic Clinical Data: An Early Readmission Risk Score," *Medical Care*, 55, 267–275. [9]

Tang, Y., Zeng, X., Feng, Y., Chen, Q., Liu, Z., Luo, H., Zha, L., and Yu, Z. (2021), "Association of Systemic Immune-Inflammation Index with Short-Term Mortality of Congestive Heart Failure: A Retrospective Cohort Study," *Frontiers in Cardiovascular Medicine*, 8, 753133. [9]

Tibshirani, R. (1996), "Regression Shrinkage and Selection Via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [1]

Van Engelen, J. E., and Hoos, H. H. (2020), "A Survey on Semi-Supervised Learning," *Machine Learning*, 109, 373–440. [2]

Verzelen, N. (2012), "Minimax Risks for Sparse Regressions: Ultra-High Dimensional Phenomenons," *Electronic Journal of Statistics*, 6, 38–90. [3]

Wasserman, L., and Lafferty, J. D. (2008), "Statistical Analysis of Semi-Supervised Regression," in *Advances in Neural Information Processing Systems* (Vol. 20), eds. J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, pp. 801–808, Curran Associates, Inc. [2]

Wei, C., Shen, K., Chen, Y., and Ma, T. (2020), "Theoretical Analysis of Self-Training with Deep Networks on Unlabeled Data," arXiv preprint arXiv:2010.03622. [2]

Weisenthal, S. J., Quill, C., Farooq, S., Kautz, H., and Zand, M. S. (2018), "Predicting Acute Kidney Injury at Hospital Re-entry Using High-Dimensional Electronic Health Record Data," *PLoS One*, 13, e0204920. [1]

Xie, Q., Luong, M.-T., Hovy, E., and Le, Q. V. (2020), "Self-Training with Noisy Student Improves Imagenet Classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10687–10698. [2]

Xue, Y., Klabjan, D., and Luo, Y. (2019), "Predicting ICU Readmission Using Grouped Physiological and Medication Trends," *Artificial Intelligence in Medicine*, 95, 27–37. [9]

Yang, X., Song, Z., King, I., and Xu, Z. (2021), "A Survey on Deep Semi-Supervised Learning," arXiv preprint arXiv:2103.00550. [2]

Zhang, A., Brown, L. D., and Cai, T. T. (2019), "Semi-Supervised Inference: General Theory and Estimation of Means," *The Annals of Statistics*, 47, 2538–2566. [2]

Zhao, J., and Leng, C. (2016), "An Analysis of Penalized Interaction Models," *Bernoulli*, 22, 1937–1961. [3,8]

Zhu, X. J. (2005), "Semi-Supervised Learning Literature Survey," Technical Report, University of Wisconsin-Madison Department of Computer Sciences. [1,2]

Supplemental Materials for “Optimal and Safe
Estimation for High-Dimensional Semi-Supervised
Learning”

October 19, 2023

Supplementary Materials

S.1 Sub-Gaussian variable and vector

To characterize the tail behavior of random variables, we introduce the following definition.

Definition S.1 (Sub-Gaussian variable and vector). A random variable X is called sub-Gaussian if there exists a positive constant K_2 such that $\mathbb{P}(|X| > t) \leq \exp(1 - t^2/K_2^2)$ for all $t \geq 0$. The sub-Gaussian norm of X is defined as $\|X\|_{\psi_2} = \sup_{p \geq 1} p^{-1/2} (\mathbb{E}|X|^p)^{1/p}$. A vector $\mathbf{X} \in \mathbb{R}^p$ is a sub-Gaussian vector if the one-dimensional marginals $\mathbf{v}^T \mathbf{X}$ are sub-Gaussian for all $\mathbf{v} \in \mathbb{R}^p$, and its sub-Gaussian norm is defined as $\|\mathbf{X}\|_{\psi_2} = \sup_{\|\mathbf{v}\|_2=1} \|\mathbf{v}^T \mathbf{X}\|_{\psi_2}$.

S.2 Illustration of Convergence Rate

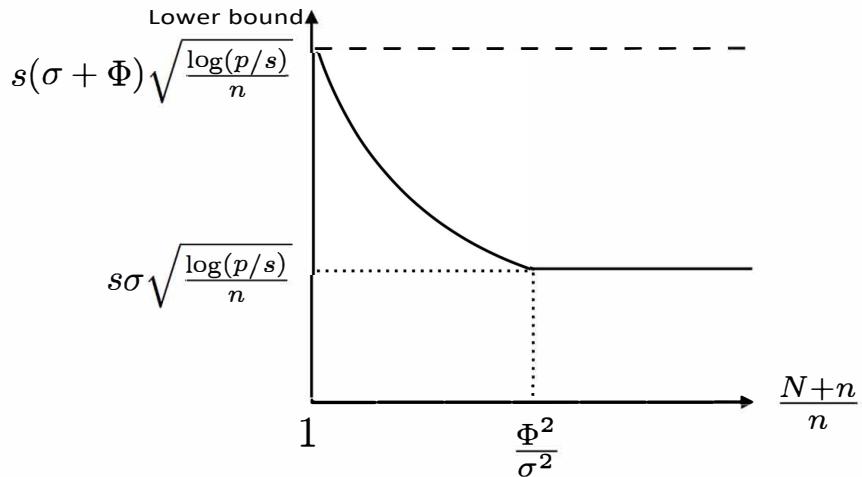


Figure 1: Plot of the lower bound with $q = 1$ in Theorem 2.2 (the solid curve) and the upper bound from Dantzig selector in (2.3) (the dashed line), against the value of $(N+n)/n$. The region between the two corresponds to the gap between the lower bound for semi-supervised estimators and the upper bound obtained from the supervised estimators. In this plot, we fix n and vary the value of N .

S.3 Proofs

Throughout the proofs without causing extra confusions, for notation simplicity, we use C, C', c_1, c_2 etc. to denote generic constants whose values can change from time to time.

S.3.1 Preliminary Lemmas

We start with several basic lemmas that we will apply in our proofs.

Lemma S.1 (Lemma B.1 in [Chernozhukov et al. \(2018\)](#)). Let $\{X_n\}, \{Y_n\}$ be sequences of random variables. If for any $c > 0$, $\mathbb{P}(|X_n| > c|Y_n) = o_p(1)$. Then $X_n = o_p(1)$.

Lemma S.2 (Nemirovski moment inequality, Lemma 14.24 in [Bhlmann and Van de Geer \(2011\)](#)). For $m \geq 1$ and $p > e^{m-1}$, we have

$$\mathbb{E} \left[\max_{1 \leq k \leq p} \left| \sum_{i=1}^n (\gamma_k(Z_i) - \mathbb{E}[\gamma_k(Z_i)]) \right|^m \right] \leq (8 \log 2p)^{\frac{m}{2}} \mathbb{E} \left[\left(\max_{1 \leq k \leq p} \sum_{i=1}^n \gamma_k^2(Z_i) \right)^{m/2} \right] \quad (\text{S.1})$$

Lemma S.3 (Theorem 3.1 in [Rudelson and Zhou \(2012\)](#)). Assume that $\mathbf{X} \in \mathbb{R}^{n \times p}$ has zero mean and covariance Σ . Furthermore, assume that the rows of $\mathbf{X}\Sigma^{-1/2} \in \mathbb{R}^{n \times p}$ are independent sub-gaussian random vector with a bounded sub-gaussian constant and $\Lambda_{\min}(\Sigma) > C_{\min} > 0$, $\max_{1 \leq j \leq p} \Sigma_{jj} = O(1)$. Set $0 < \delta < 1$, $0 < s_0 < p$, and $L > 0$. Define the following event,

$$\mathcal{B}_\delta(n, s_0, L) = \{ \mathbf{X} \in \mathbb{R}^{n \times p} : (1 - \delta)\sqrt{C_{\min}} \leq \frac{\|\mathbf{X}\mathbf{v}\|_2}{\sqrt{n}\|\mathbf{v}\|_2}, \forall \mathbf{v} \in \mathcal{C}(s_0, L) \text{ s.t. } \mathbf{v} \neq 0 \}. \quad (\text{S.2})$$

and $\mathcal{C}(s_0, L) = \{ \boldsymbol{\theta} \in \mathbb{R}^p : \exists S \subseteq \{1, \dots, p\}, |S| = s_0, \|\boldsymbol{\theta}_{S^c}\|_1 \leq L\|\boldsymbol{\theta}_S\|_1 \}$. Then, there exists a constant $c_1 = c(L, \delta)$ such that, for sample size $n \geq c_1 s_0 \log(p/s_0)$, we have

$$\mathbb{P}(\mathcal{B}_\delta(n, s_0, L)) \geq 1 - e^{-\delta^2 n}. \quad (\text{S.3})$$

Lemma S.4. Assume that Assumption 3.1 holds. Consider the Dantzig selector $\widehat{\boldsymbol{\theta}}_D$ in (2.2)

with $\lambda_D \asymp K_1 \sqrt{\frac{(\sigma^2 + \Phi^2) \log p}{n}}$. We have

$$\|\widehat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*\|_1 = O_p(s\lambda_D), \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n [X_i^T(\widehat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*)]^2 = O_p(s\lambda_D^2). \quad (\text{S.4})$$

Moreover, we have

$$\|\widehat{\boldsymbol{\Sigma}}_n(\widehat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*)\|_\infty = O_p(\lambda_D),$$

where $\widehat{\boldsymbol{\Sigma}}_n = \frac{1}{n} \sum_{i=1}^n X_i^{\otimes 2}$.

Proof. The proof of the convergence rate of $\widehat{\boldsymbol{\theta}}_D$ in (S.4) is similar to Theorem 7.1 in [Bickel et al. \(2009\)](#). The key step is to derive

$$\left\| \frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*) \right\|_\infty \lesssim K_1(\sigma^2 + \Phi^2)^{1/2} \sqrt{\frac{\log p}{n}},$$

which is implied by Lemma S.2 together with $\mathbb{E}(Y_i - X_i^T \boldsymbol{\theta}^*)^2 = \sigma^2 + \Phi^2$ and $\|X_i\|_\infty \leq K_1$.

The rest of the proof is omitted. To show the rate of $\|\widehat{\boldsymbol{\Sigma}}_n(\widehat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*)\|_\infty$, we note that, with $\lambda_D = CK_1 \sqrt{\frac{(\sigma^2 + \Phi^2) \log p}{n}}$ for some sufficiently large C , we have

$$\begin{aligned} \|\widehat{\boldsymbol{\Sigma}}_n(\widehat{\boldsymbol{\theta}}_D - \boldsymbol{\theta}^*)\|_\infty &\leq \left\| \frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \widehat{\boldsymbol{\theta}}_D) \right\|_\infty + \left\| \frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*) \right\|_\infty \\ &\leq 2\lambda_D, \end{aligned}$$

where we invoke the constraint of $\widehat{\boldsymbol{\theta}}_D$ as the definition of Dantzig selector in the last step. \square

S.3.2 Proof of Theorem 2.2

To simplify the notation, we use $\mathcal{P}_{\Phi, \sigma}$ to denote $\mathcal{P}_{\Phi, \sigma}^{\text{pairwise}}$ or $\mathcal{P}_{\Phi, \sigma}^{\text{additive}}$ in the proof. The proof is very similar under these two classes, and we will point out the difference when it occurs.

The lower bound consists of two parts. First, by taking $f(X) = X^T \boldsymbol{\theta}^*$, that is when the linear model is correctly specified, the proof in Proposition 6.4 of [Verzelen \(2012\)](#) with

conditions $s \log(p/s) \leq C'n$ and $2 \leq s \leq (n-1)/4$ directly implies

$$\inf_{\hat{\theta}} \sup_{P_{X,Y} \in \mathcal{P}_{\Phi,\sigma}} \mathbb{P}_{P_{X,Y}} \left[\|\hat{\theta} - \theta^*\|_q \gtrsim c_1 s^{1/q} \sigma \sqrt{\frac{\log(p/s)}{n}} \right] > c_2.$$

We note that taking $f(X) = X^T \theta^*$ is feasible under $\mathcal{P}_{\Phi,\sigma}^{\text{pairwise}}$ and $\mathcal{P}_{\Phi,\sigma}^{\text{additive}}$, and we also use the assumption $\bar{s} \geq s$.

Second, to establish the lower bound $s^{1/q} \Phi \sqrt{\log(p/s)/(n+N)}$, we first construct a set of hypotheses and then apply Theorem 2.7 in [Tsybakov \(2008\)](#). In this case, we take $f(X) = b f_u(X_u)$, for some $u \subseteq \{1, 2, \dots, p\}$ with $a = |u| \leq s/2$ and $b \in \mathbb{R}$ to be set later. For example, under $\mathcal{P}_{\Phi,\sigma}^{\text{pairwise}}$, we can set $f_u(X_u) = X_1 X_2$; under $\mathcal{P}_{\Phi,\sigma}^{\text{additive}}$, we can set $f_u(X_u) = X_1^2$ (note that x^2 is a second order smooth function). Without the loss of generality, we set $u = \{1, \dots, a\}$.

Define the set $\mathcal{M} = \{x \in \{0, 1\}^{p-a} : \|x\|_0 = s/2\}$. It follows from the Varshamov-Gilbert bound (e.g., Lemma 2.9 in [Tsybakov \(2008\)](#)) and Lemma A.3 in [Rigollet and Tsybakov \(2011\)](#) that there exists a subset \mathcal{M}' of \mathcal{M} such that for any x, x' in \mathcal{M}' with $x \neq x'$, we have

$$\rho_H(x, x') > \frac{s}{16}, \quad \text{and} \quad \log |\mathcal{M}'| \geq c'_1 s \log \left(\frac{p-a}{s} \right), \quad (\text{S.5})$$

where ρ_H denotes the Hamming distance and $c'_1 > 0$ is an absolute constant. Denote the element of the finite set \mathcal{M}' by \mathbf{w}^j and the index set for the nonzero entries by $[j]$. For each $\mathbf{w}^j \in \mathbb{R}^{p-a}$, we extend the vector by adding zero at the first a entries, i.e. the corresponding element in u to obtain a p -dimensional vector $(0_a, \mathbf{w}^j)$, and for notational simplicity, we still call it \mathbf{w}^j .

Next, we construct a finite set of hypotheses by perturbing the distribution of X . Denote $N = (N_1, \dots, N_p) \sim \mathcal{N}(0, I_p)$. Let us consider the following hypothesis

$$H_0 : X = (X_1, \dots, X_p) \text{ with } X_\ell = N_\ell,$$

and under H_0 , denote $\theta^* = \arg \min_{\theta \in \mathbb{R}^p} \mathbb{E}_0[(f(X) - X^T \theta)^2]$. It is easily seen that $\theta_{u^c}^* = 0$.

With $f^\perp(X) = f(X) - X^T \boldsymbol{\theta}^*$, we can choose b so that $\mathbb{E}_0[(f^\perp(X))^2] = \Phi^2$. We construct other hypotheses as

$$H_j : X = (X_1, \dots, X_p) \text{ with } X_\ell = \begin{cases} \frac{\rho}{\Phi} f^\perp(N) + \sqrt{1 - \rho^2} N_\ell, & \text{if } \mathbf{w}_\ell^j = 1, \\ N_\ell, & \text{if } \mathbf{w}_\ell^j = 0, \end{cases}$$

for $j = 1, \dots, |\mathcal{M}'|$, where $\rho > 0$ is a quantity to be chosen later. Let \mathbb{E}_j denote the expectation under H_j . Clearly, $\mathbb{E}_j(X_\ell) = 0$. After some simple calculation, we can verify that for $j = 1, \dots, |\mathcal{M}'|$,

$$\mathbb{E}_j[X_\ell X_m] = \begin{cases} \rho^2 & \text{if } \mathbf{w}_\ell^j \text{ and } \mathbf{w}_m^j \neq 0 \text{ and } \ell \neq m, \\ 1 & \text{if } \ell = m, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{S.6})$$

Denote by $\mathbf{M} = \mathbb{E}_j[XX^T]$ the covariance matrix of X under H_j . From (S.6), there exists a permutation matrix \mathbf{P} such that $\mathbf{M} = \mathbf{P} \mathbf{B} \mathbf{P}^T$, where $\mathbf{B} = \text{diag}(\mathbf{A}, I_{p-\frac{s}{2}})$ and \mathbf{A} is a $s/2$ -dimensional equicorrelation matrix with the off-diagonal entry ρ^2 and the diagonal entry 1. Assume that the (λ, \mathbf{v}) are the eigenvalue and corresponding eigenvector of \mathbf{B} . Following the definition of eigenvalues, $(\lambda, \mathbf{P}\mathbf{v})$ are the eigenvalue and eigenvector of \mathbf{M} . As a result, $\lambda_{\min}(\mathbf{M}) = \lambda_{\min}(\mathbf{B}) = \lambda_{\min}(\mathbf{A}) = 1 - \rho^2$, where the last two equalities follow from the property of the block diagonal matrix and equicorrelation matrix. With the choice of ρ as specified in (S.11), we can derive that $\lambda_{\min}(\mathbf{M}) > 1/2$ and therefore \mathbf{M} is positive definite.

Our next step is to verify that the corresponding estimand $\boldsymbol{\theta}^j = \{\mathbb{E}_j(XX^T)\}^{-1} \mathbb{E}_j(Xf(X))$ is s -sparse. To this end, we first note that

$$\mathbb{E}_j[X_\ell f(X)] = \begin{cases} \rho \Phi, & \text{if } \mathbf{w}_\ell^j = 1, \\ \mathbb{E}_j[N_\ell f(N)] = \boldsymbol{\theta}_\ell^*, & \text{if } \mathbf{w}_\ell^j = 0. \end{cases}$$

To calculate \mathbf{M}^{-1} , we first rewrite the equicorrelation matrix \mathbf{A} as $\mathbf{A} = (1 - \rho^2)I_{\frac{s}{2}} + \rho^2 \mathbf{1}_{\frac{s}{2}} \mathbf{1}_{\frac{s}{2}}^T$,

where $\mathbf{1}_{\frac{s}{2}}$ is a $\frac{s}{2}$ -dimensional vector of 1. The Woodbury formula implies

$$\mathbf{A}^{-1} = \frac{1}{1 - \rho^2} I_{\frac{s}{2}} - \frac{\rho^2}{(1 - \rho^2)(1 + (\frac{s}{2} - 1)\rho^2)} \mathbf{1}_{\frac{s}{2}} \mathbf{1}_{\frac{s}{2}}^T.$$

Thus, we obtain $\mathbf{M}^{-1} = \mathbf{P}\mathbf{B}^{-1}\mathbf{P}^T$ with $\mathbf{B}^{-1} = \text{diag}(\mathbf{A}^{-1}, I_{p-\frac{s}{2}})$. Let \mathbf{e}_t denote the canonical basis in \mathbb{R}^p with the t -th entry being 1 and the rest being 0. Note that the permutation matrix \mathbf{P} can be written as $\mathbf{P} = (\mathbf{e}_{j_1}, \dots, \mathbf{e}_{j_{s/2}}, \mathbf{e}_{j_{s/2+1}}, \dots, \mathbf{e}_{j_p})$, where the indexes $j_1, \dots, j_{s/2}$ belong to $[j]$ and $j_{s/2+1}, \dots, j_p$ are not in $[j]$. Combining the above argument, the estimand $\boldsymbol{\theta}^j$ under hypothesis H_j is given by

$$\begin{aligned} \boldsymbol{\theta}^j &= \mathbf{P}\mathbf{B}^{-1}\mathbf{P}^T \mathbb{E}_j[Xf(X)] = \mathbf{P}\mathbf{B}^{-1} \begin{pmatrix} \rho\Phi \mathbf{1}_{s/2} \\ \boldsymbol{\theta}_{-[j]}^* \end{pmatrix} = \mathbf{P} \begin{pmatrix} \rho\Phi \mathbf{A}^{-1} \mathbf{1}_s \\ \boldsymbol{\theta}_{-[j]}^* \end{pmatrix} \\ &= \begin{cases} \frac{\rho\Phi}{1 + (\frac{s}{2} - 1)\rho^2}, & \text{if } \mathbf{w}_\ell^j = 1, \\ \boldsymbol{\theta}_l^*, & \text{if } \mathbf{w}_\ell^j = 0. \end{cases}, \end{aligned} \quad (\text{S.7})$$

from which we know $\boldsymbol{\theta}^j$ is s -sparse since $\|\boldsymbol{\theta}^j\|_0 \leq \frac{s}{2}$.

In the sequel, we will verify $\mathbb{E}_j(f(X) - X^T \boldsymbol{\theta}^j)^2 \leq \Phi^2$ holds. Recall that $N = (N_1, \dots, N_p) \sim \mathcal{N}(0, I_p)$. By (S.7), we have

$$\begin{aligned} \mathbb{E}_j(f(X) - X^T \boldsymbol{\theta}^j)^2 &= \mathbb{E}_0 \left[(f(N) - N^T \boldsymbol{\theta}^* - \frac{s\rho^2 f^\perp(N)}{2(1 + (s/2 - 1)\rho^2)} - \frac{\rho\Phi\sqrt{1 - \rho^2}}{1 + (s/2 - 1)\rho^2} \sum_{k \in [j]} N_k)^2 \right] \\ &= (1 - \frac{s\rho^2}{2(1 + (s/2 - 1)\rho^2)})^2 \Phi^2 + \frac{s\rho^2 \Phi^2 (1 - \rho^2)}{2(1 + (s/2 - 1)\rho^2)^2} \\ &= \Phi^2 \frac{(1 - \rho^2)^2 + s\rho^2(1 - \rho^2)/2}{(1 + (s/2 - 1)\rho^2)^2} \\ &= \Phi^2 \frac{1 - (s/2 - 1)\rho^4 + (s/2 - 2)\rho^2}{1 + (s - 2)\rho^2 + (s/2 - 1)^2\rho^4} < \Phi^2 \end{aligned}$$

Due to the choice of b , under H_0 we have $\mathbb{E}_0(f(X) - X^T \boldsymbol{\theta}^0)^2 = \Phi^2$. Therefore, we have shown that the distribution of (X, Y) under the hypotheses H_j for $j = 0, \dots, |\mathcal{M}'|$ belongs to

the class of distributions $\mathcal{P}_{\Phi,\sigma}$.

To apply Theorem 2.7 in [Tsybakov \(2008\)](#), we need to (1) lower bound $\|\boldsymbol{\theta}^j - \boldsymbol{\theta}^{j'}\|_q$ for $0 \leq j < j' \leq |\mathcal{M}'|$ and (2) upper bound the Kullback-Leibler divergence between the probability measure of the data denoted by \mathcal{P}_j and \mathcal{P}_0 under H_j and H_0 . For (1), we have from (S.7) that for $1 \leq j < j' \leq |\mathcal{M}'|$

$$\|\boldsymbol{\theta}^j - \boldsymbol{\theta}^{j'}\|_q = \frac{\Phi\rho}{1 + (s/2 - 1)\rho^2} \rho_H^{1/q}(\mathbf{w}^j, \mathbf{w}^{j'}) \geq \frac{s^{1/q}\Phi\rho}{16^{1/q}(1 + (s/2 - 1)\rho^2)}, \quad (\text{S.8})$$

where the last step follows from (S.5). For $j = 0$ and $j' \geq 1$ we have

$$\|\boldsymbol{\theta}^j - \boldsymbol{\theta}^{j'}\|_q = \frac{\Phi\rho s^{1/q}}{2^{1/q}(1 + (s/2 - 1)\rho^2)}. \quad (\text{S.9})$$

To quantify the Kullback-Leibler divergence, recall that the data in matrix form can be written as $(\mathbf{Y}, \widetilde{\mathbf{X}})$, where $\widetilde{\mathbf{X}} = (X_1, \dots, X_{n+N})^T$ and $\mathbf{Y} = (Y_1, \dots, Y_n)^T$. With a slight change of notation, we use $X_{i\ell}$ to denote the ℓ -th component of X_i for $1 \leq \ell \leq p$ and $1 \leq i \leq n+N$. Under H_j , the data distribution can be decomposed as

$$\begin{aligned} p_j(\mathbf{Y}, \widetilde{\mathbf{X}}) &= p(\mathbf{Y}|\mathbf{X})p_j(\widetilde{\mathbf{X}}) = p(\mathbf{Y}|\mathbf{X}) \prod_{i=1}^{n+N} p_j(X_i) \\ &= p(\mathbf{Y}|\mathbf{X}) \prod_{i=1}^{n+N} p_j(X_{i,-\{[j],u\}})p_j(X_{i,[j]}|X_{i,u})p_j(X_{i,u}), \end{aligned}$$

where we note that the p.d.f. $p(\mathbf{Y}|\mathbf{X})$ remains the same across j and $X_{i,-\{[j],u\}}$ stands for the subvector of X_i by excluding the indexes in $\{[j], u\}$. From the above decomposition, the Kullback-Leibler divergence is given by

$$\mathcal{K}(\mathcal{P}_j, \mathcal{P}_0) = \mathbb{E}_j \left[\log \frac{p(\mathbf{Y}|\mathbf{X})p_j(\widetilde{\mathbf{X}})}{p(\mathbf{Y}|\mathbf{X})p_0(\widetilde{\mathbf{X}})} \right] = (n+N)\mathbb{E}_j \left[\log \frac{p_j(X_{i,[j]}|X_{i,u})}{p_0(X_{i,[j]}|X_{i,u})} \right].$$

Furthermore, notice that

$$\begin{aligned} p_j(X_{i,[j]}|X_{i,u}) &= \prod_{k \in [j]} p_j(X_{ik}|X_{i,u}) \\ &= \prod_{k \in [j]} \left(\frac{1}{\sqrt{2\pi(1-\rho^2)}} \exp \left\{ -\frac{(X_{ik} - \frac{\rho}{\Phi} f^\perp(X_u))^2}{2(1-\rho^2)} \right\} \right), \end{aligned}$$

$$p_0(X_{i,[j]}|X_{i,u}) = \prod_{k \in [j]} p_0(X_{ik}|X_{i,u}) = \prod_{k \in [j]} \left(\frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{X_{ik}^2}{2} \right\} \right).$$

Hence, with some index $k \in [j]$ we obtain

$$\begin{aligned} \mathcal{K}(\mathcal{P}_j, \mathcal{P}_0) &= \frac{(n+N)s}{2} \left(\mathbb{E}_j \left[-\frac{(X_{ik} - \frac{\rho}{\Phi} f^\perp(X_u))^2}{(1-\rho^2)} \right] - \log(1-\rho^2) + \mathbb{E}_j[X_{ik}^2] \right) \\ &= \frac{(n+N)s}{2} \left(1 - \log(1-\rho^2) \right. \\ &\quad \left. - \frac{1}{(1-\rho^2)} \mathbb{E}_j \left[X_{ik}^2 + \frac{\rho^2}{\Phi^2} ((f^\perp(X_u))^2 - 2\frac{\rho}{\Phi} X_{ik} f^\perp(X_u)) \right] \right) \\ &= \frac{(n+N)s}{2} \left(1 - \log(1-\rho^2) - \frac{1+\rho^2-2\rho^2}{1-\rho^2} \right) \\ &= \frac{(n+N)s}{2} \log \left(1 + \frac{\rho^2}{1-\rho^2} \right) \leq \frac{s(n+N)\rho^2}{2(1-\rho^2)}. \end{aligned} \tag{S.10}$$

We set

$$\rho = \frac{1}{4} \sqrt{\frac{c'_1 \log(p/s)}{n+N}}, \tag{S.11}$$

where c'_1 is specified in (S.5). Given the condition $\frac{(s/2-1)c'_1 \log(p/s)}{(n+N)} \leq 1$ and $s \geq 3$, we have $1-\rho^2 \geq \frac{1}{2}$. Then from (S.10) the Kullback-Leibler divergence can be bounded as follows

$$\mathcal{K}(\mathcal{P}_j, \mathcal{P}_0) \leq \frac{s(n+N)\rho^2}{2(1-\rho^2)} \leq \frac{c'_1 s \log(p/s)}{16} \leq \frac{1}{16} \log |\mathcal{M}'|,$$

where the last inequality follows from (S.5). Finally, with the choice of ρ in (S.11), we obtain from (S.8) and (S.9) that

$$\|\boldsymbol{\theta}^j - \boldsymbol{\theta}^{j'}\|_q \geq C s^{1/q} \Phi \sqrt{\frac{\log(p/s)}{n+N}}.$$

where we use the inequality $1 + (s/2 - 1)\rho^2 \leq 17/16$ and C is a constant. We complete the proof by applying Theorem 2.7 in [Tsybakov \(2008\)](#).

S.3.3 Proof of Theorem 3.2

Proof. Here, we show a general version with surrogate function $h(X)$. First, we can rewrite that

$$\begin{aligned}\widehat{\xi}_j &= \sum_{i \in D_j^*} \frac{X_i Y_i - \widehat{h}^{-j}(X_i) X_i}{n_j} - \sum_{i \in D_j} \frac{\widehat{h}^{-j}(X_i) X_i}{n_j + N_j} \\ &= \sum_{i \in D_j^*} \frac{X_i Y_i - f(X_i) X_i}{n_j} - \sum_{i \in D_j} \frac{f(X_i) X_i}{n_j + N_j} + \sum_{i \in D_j^*} \frac{(f(X_i) - h(X_i)) X_i}{n_j} \\ &\quad - \sum_{i \in D_j} \frac{(f(X_i) - h(X_i)) X_i}{n_j + N_j} + \sum_{i \in D_j^*} \frac{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i}{n_j} - \sum_{i \in D_j} \frac{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i}{n_j + N_j}.\end{aligned}$$

Denote

$$I_1 = \sum_{i \in D_j^*} \frac{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i}{n_j} - \sum_{i \in D_j} \frac{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i}{n_j + N_j}. \quad (\text{S.12})$$

$$I_2 = \sum_{i \in D_j^*} \frac{(f(X_i) - h(X_i)) X_i}{n_j} - \sum_{i \in D_j} \frac{(f(X_i) - h(X_i)) X_i}{n_j + N_j}. \quad (\text{S.13})$$

Next, we aim to show that $\|I_1\|_\infty = O_p(K_1 b_n \sqrt{\frac{\log p}{n}})$. To this end, we further decompose I as

$$\begin{aligned}\|I_1\|_\infty &\leq \left\| \frac{1}{n_j} \sum_{i \in D_j^*} \{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i\} - \mathbb{E}_{D_{-j}^*} [(h(X) - \widehat{h}^{-j}(X)) X] \right\|_\infty \\ &\quad + \left\| \frac{1}{n_j + N_j} \sum_{i \in D_j} \{(h(X_i) - \widehat{h}^{-j}(X_i)) X_i\} - \mathbb{E}_{D_{-j}^*} [(h(X) - \widehat{h}^{-j}(X)) X] \right\|_\infty, \quad (\text{S.14})\end{aligned}$$

where $\mathbb{E}_{D_{-j}^*}$ denotes the conditional expectation given the data in $D_{-j}^* = D^* \setminus D_j^*$. Let us denote $g_k(X) = (h(X) - \widehat{h}^{-j}(X)) X_k$ and $\gamma_k(X) = g_k(X) - \mathbb{E}_{D_{-j}^*}[g_k(X)]$. From (S.14), it suffices to upper bound $\max_{1 \leq k \leq p} \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i)$ and $\max_{1 \leq k \leq p} \frac{1}{n_j + N_j} \sum_{i \in D_j} \gamma_k(X_i)$, respectively.

We know

$$\mathbb{E}_{D_{-j}^*} \left[\max_{1 \leq k \leq p} \sum_{i \in D_j^*} g_k^2(X_i) \right] \leq n_j K_1^2 \|\hat{h}^{-j} - h\|_2^2, \quad (\text{S.15})$$

which follows from $\|X\|_\infty \leq K_1$ in Assumption 3.1. Therefore, with the application of lemma S.2 by choosing $m = 2$, we can show

$$\mathbb{E}_{D_{-j}^*} \left[\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right|^2 \right] \leq K_1^2 \frac{8 \log(2p)}{n_j} \|\hat{h}^{-j} - h\|_2^2.$$

Furthermore, the Markov inequality implies for any $c > 0$

$$\begin{aligned} & \mathbb{P} \left(\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right| \geq c K_1 b_n \sqrt{\frac{\log 2p}{n_j}} \middle| D_{-j}^* \right) \\ & \leq \left(\frac{\mathbb{E}_{D_{-j}^*} [\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right|^2]}{c^2 K_1^2 b_n^2} \frac{n_j}{\log(2p)} \right) \wedge 1 \\ & \leq \left(\frac{8 \|\hat{h}^{-j} - h\|_2^2}{c^2 b_n^2} \right) \wedge 1. \end{aligned} \quad (\text{S.16})$$

For any $\epsilon > 0$, let c' be a sufficiently large constant such that the event $\mathcal{E} = \{\|\hat{h}^{-j} - h\|_2^2 \leq c' b_n^2\}$ holds with probability at least $1 - \epsilon$. From (S.16), we know that

$$\begin{aligned} & \mathbb{P} \left(\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right| \geq c K_1 b_n \sqrt{\frac{\log 2p}{n_j}} \right) \\ & = \mathbb{E} \left[\mathbb{P} \left(\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right| \geq c K_1 b_n \sqrt{\frac{\log 2p}{n_j}} \middle| D_{-j}^* \right) \right] \\ & \leq \mathbb{E} \left[\left(\frac{8 \|\hat{h}^{-j} - h\|_2^2}{c^2 b_n^2} \wedge 1 \right) I(\mathcal{E}) \right] + \mathbb{E} \left[\left(\frac{8 \|\hat{h}^{-j} - h\|_2^2}{c^2 b_n^2} \wedge 1 \right) I(\mathcal{E}^c) \right] \\ & \leq \frac{8c'}{c^2} + \mathbb{P}(\mathcal{E}^c) \leq 2\epsilon. \end{aligned}$$

where the last step holds by taking $c^2 = 8c'/\epsilon$ and the definition $\mathbb{P}(\mathcal{E}^c) \leq \epsilon$. This implies

$$\max_{1 \leq k \leq p} \left| \frac{1}{n_j} \sum_{i \in D_j^*} \gamma_k(X_i) \right| = O_p(K_1 b_n \sqrt{\frac{\log p}{n_j}}).$$

Following the same argument, the following probability bound holds

$$\max_{1 \leq k \leq p} \left| \frac{1}{n_j + N_j} \sum_{i \in D_j} \gamma_k(X_i) \right| = O_p(K_1 b_n \sqrt{\frac{\log p}{n_j + N_j}}).$$

What is more, if we denote $\|f - h\|_2^2 = \mathbb{E}[(f(X) - h(X))^2]$, we can also derive

$$\|I_2\|_\infty = O_p(K_1 \|f - h\|_2 \sqrt{\frac{\log p}{n_j}}).$$

The rest of the proof follows the same line as in the proof of Theorem 7.1 in [Bickel et al. \(2009\)](#). Recall that we assume for $j = \{1, 2\}$, $n_j = n/2$ and $N_j = N/2$. We can show that

$$\begin{aligned} & \left\| \frac{\sum_{i=1}^{n+N} X_i X_i^T}{(n+N)} \boldsymbol{\theta}^* - \frac{\widehat{\boldsymbol{\xi}}_1 + \widehat{\boldsymbol{\xi}}_2}{2} \right\|_\infty \\ &= \left\| \frac{\sum_{i=1}^{n+N} X_i (X_i^T \boldsymbol{\theta}^* - f(X_i))}{(n+N)} + \frac{\sum_{i=1}^n X_i^T (Y_i - f(X_i))}{n} \right\|_\infty + O_p(K_1 b_n \sqrt{\frac{\log p}{n}}) \\ & \quad + O_p(K_1 \|f - h\|_2 \sqrt{\frac{\log p}{n_j}}) \\ &\leq \left\| \frac{\sum_{i=1}^{n+N} X_i (X_i^T \boldsymbol{\theta}^* - f(X_i))}{(n+N)} \right\|_\infty + \left\| \frac{\sum_{i=1}^n X_i^T (Y_i - f(X_i))}{n} \right\|_\infty + O_p(K_1 b_n \sqrt{\frac{\log p}{n}}) \\ & \quad + O_p(K_1 \|f - h\|_2 \sqrt{\frac{\log p}{n_j}}) \\ &= O_p(K_1 \Phi \sqrt{\frac{\log p}{n+N}}) + O_p(K_1 \sigma \sqrt{\frac{\log p}{n}}) + O_p(K_1 b_n \sqrt{\frac{\log p}{n}}) + O_p(K_1 \|f - h\|_2 \sqrt{\frac{\log p}{n_j}}). \end{aligned} \tag{S.17}$$

where the last probability bound holds by the same argument in the proof of example 14.3 in [Bhlmann and Van de Geer \(2011\)](#).

Then, if we set $\lambda_{SD} = C' K_1 (\Phi \sqrt{\frac{\log p}{n+N}} + (\sigma + b_n + \|f - h\|_2) \sqrt{\frac{\log p}{n}})$ for sufficiently large C' , (S.17) implies

$$\lambda_{SD} \geq \left\| \frac{\sum_{i=1}^{n+N} X_i X_i^T \boldsymbol{\theta}^* - \frac{\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2}{2}}{(n+N)} \right\|_\infty \quad (\text{S.18})$$

holds with probability tending to 1. Let $\boldsymbol{\delta} = \hat{\boldsymbol{\theta}}_{SD} - \boldsymbol{\theta}^*$. By the construction of Dantzig estimator, when (S.18) holds, we have $\|\boldsymbol{\delta}_{T^c}\|_1 \leq \|\boldsymbol{\delta}_T\|_1$ where T denotes the support of $\boldsymbol{\theta}^*$ and

$$\frac{1}{n+N} \|\widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} \boldsymbol{\delta}\|_\infty \leq \left\| \frac{\widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} \hat{\boldsymbol{\theta}}_{SD}}{n+N} - \frac{\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2}{2} \right\|_\infty + \left\| \frac{\widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} \boldsymbol{\theta}^*}{n+N} - \frac{\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2}{2} \right\|_\infty \leq 2\lambda_{SD}.$$

Therefore,

$$\begin{aligned} \frac{1}{n+N} \|\widetilde{\mathbf{X}} \boldsymbol{\delta}\|_2^2 &= \frac{1}{n+N} \boldsymbol{\delta}^T \widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}} \boldsymbol{\delta} \\ &\leq \frac{1}{n+N} \|\boldsymbol{\delta}^T \widetilde{\mathbf{X}}^T \widetilde{\mathbf{X}}\|_\infty \|\boldsymbol{\delta}\|_1 \\ &\leq 2\lambda_{SD} \times 2\|\boldsymbol{\delta}_T\|_1 \leq 4\lambda_{SD} \sqrt{s} \|\boldsymbol{\delta}_T\|_2. \end{aligned}$$

With smallest eigenvalue condition (A1) and Lemma S.3, we know on the event of $\mathcal{B}_{c_1}(n+N, s, 1)$, $\frac{1}{n+N} \|\widetilde{\mathbf{X}} \boldsymbol{\delta}\|_2^2 \geq (1 - c_1)^2 C_{\min} \|\boldsymbol{\delta}_T\|_2^2$. Therefore,

$$\|\boldsymbol{\delta}_T\|_2 \leq \frac{4\lambda_{SD} \sqrt{s}}{(1 - c_1)^2 C_{\min}}.$$

Above all, we know $\|\boldsymbol{\delta}\|_1 \leq 2\|\boldsymbol{\delta}_T\|_1 \leq 2\sqrt{s} \|\boldsymbol{\delta}_T\|_2 \leq \frac{8\lambda_{SD} s}{(1 - c_1)^2 C_{\min}}$. As a byproduct, we can show that $\|\boldsymbol{\delta}\|_2 \lesssim \lambda_D \sqrt{s}$. This completes the proof. \square

S.3.4 Proof of Theorem 4.1

Before the proof starts, we need an additional notation for the largest restricted eigenvalue.

Denote

$$\phi_{\max}(L) = \max_{1 \leq \|v\|_0 \leq L} \frac{v^T \widehat{\boldsymbol{\Sigma}}_{n+N} v}{\|v\|_2^2}.$$

Then we know $\phi_{\max}(\cdot)$ is an increasing function. We further use the same notation to denote the event for the upper bound parallel to the lower bound in definition (S.2)

$$\mathcal{B}_\delta^U(n, s_0, L) = \{\mathbf{X} \in \mathbb{R}^{n \times p} : \frac{\|\mathbf{X}\mathbf{v}\|_2}{\sqrt{n}\|\mathbf{v}\|_2} \leq (1 + \delta)\sqrt{C_{\max}}, \forall \mathbf{v} \in \mathcal{C}(s_0, L) \text{ s.t. } \mathbf{v} \neq 0\}, \quad (\text{S.19})$$

and under same condition in Lemma S.3,

$$\mathbb{P}(\mathcal{B}_\delta^U(n, s_0, L)) \geq 1 - e^{-\delta^2 n}. \quad (\text{S.20})$$

Lemma 3.5 in [Javanmard and Montanari \(2018\)](#) shows that $|\widehat{T}_1| \lesssim s$ with high probability under the assumption that $s(\log p)^2 = o(n)$. With those notations above, we can show the following lemma ensuring the sparsity of \widehat{T}_2 in a similar spirit.

Lemma S.5. Suppose Assumption 3.1 holds, $\Lambda_{\max}(\boldsymbol{\Sigma}) \leq C_{\max} < \infty$ and $s(\log p)^2 = o(n + N)$. On the event of $\mathcal{B}_\delta(n, s, 3) \cap \mathcal{B}_\delta^U(n, s, 3) \cap \{\lambda_{SL} \geq 2\|\boldsymbol{\Sigma}_{n+N}\boldsymbol{\theta}^* - \widehat{\boldsymbol{\xi}}\|_\infty\}$, we can show

$$|\widehat{T}_2| \leq \frac{64C_{\max}s}{(1 - c_1)^2 C_{\min}}.$$

Proof. Recall the estimator $\widehat{\boldsymbol{\theta}}_{SL}$ is defined as

$$\widehat{\boldsymbol{\theta}}_{SL} = \underset{\boldsymbol{\theta} \in \mathbb{R}^p}{\operatorname{argmin}} \sum_{j=1}^2 \frac{\sum_{i \in D_j} (\widehat{h}^{-j}(X_i) - X_i^T \boldsymbol{\theta})^2}{n + N} - \frac{2 \sum_{i \in D_j^*} (Y_i - \widehat{h}^{-j}(X_i)) X_i^T \boldsymbol{\theta}}{n} + 2\lambda_{SL}\|\boldsymbol{\theta}\|_1. \quad (\text{S.21})$$

By KKT condition, we have

$$\sum_{j=1}^2 \frac{\sum_{i \in D_j} (\widehat{h}^{-j}(X_i) - X_i^T \widehat{\boldsymbol{\theta}}_{SL}) X_i}{n + N} + \frac{\sum_{i \in D_j^*} (Y_i - \widehat{h}^{-j}(X_i)) X_i}{n} = \lambda_{SL} v(\widehat{\boldsymbol{\theta}}_{SL}), \quad v(\boldsymbol{\theta}) \in \partial\|\boldsymbol{\theta}\|_1.$$

Hence we know

$$\begin{aligned} \frac{\sum_{i=1}^{n+N} (X_i^T \boldsymbol{\theta}^* - X_i^T \widehat{\boldsymbol{\theta}}_{SL}) X_i}{n+N} &= \lambda_{SL} v(\widehat{\boldsymbol{\theta}}_{SL}^{(j)}) \\ &- \sum_{j=1}^2 \left(\frac{\sum_{i \in D_j^*} (Y_i - \widehat{h}^{-j}(X_i)) X_i}{n} + \frac{\sum_{i \in D_j} (\widehat{h}^{-j}(X_i) - \widehat{X}_i^T \boldsymbol{\theta}^*) X_i}{n+N} \right), \end{aligned}$$

and by previous reasoning in Theorem 3.2, we know

$$\left\| \sum_{j=1}^2 \left(\frac{\sum_{i \in D_j^*} (Y_i - \widehat{h}^{-j}(X_i)) X_i}{n} + \frac{\sum_{i \in D_j} (\widehat{h}^{-j}(X_i) - \widehat{X}_i^T \boldsymbol{\theta}^*) X_i}{n+N} \right) \right\|_\infty \leq \frac{\lambda_{SL}}{2}.$$

Therefore, for $k \in \widehat{T}_2$, $\left| \frac{\sum_{i \in D_j} (X_i^T \boldsymbol{\theta}^* - X_i^T \widehat{\boldsymbol{\theta}}_{SL}) X_i^k}{n+N} \right| \geq \lambda_{SL}/2$. We square both sides of the inequality and sum over $k \in \widehat{T}_2$. With $\Delta = \widetilde{\mathbf{X}}(\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}_{SL})$, we attain that

$$\frac{\lambda_{SL}^2}{4} |\widehat{T}_2| \leq \frac{\sum_{k \in \widehat{T}_2} (\Delta^T \widetilde{\mathbf{X}}_{\cdot k})^2}{n+N} \leq \|(\widehat{\boldsymbol{\Sigma}}_{n+N})_{\widehat{T}_2}\|_2^2 \|\Delta\|_2^2. \quad (\text{S.22})$$

First, by a standard argument for lasso estimator, we can show the same result as for the dantzig selector in Theorem 3.2 and we know $\|\Delta\|_2^2 \leq 16s\lambda_{SL}^2 / ((1-c_1)^2 C_{\min})$ on the event of $\mathcal{B}_\delta(n, s, 3) \cap \{\lambda_{SL} \geq 2\|\widehat{\boldsymbol{\Sigma}}_{n+N}\boldsymbol{\theta}^* - \widehat{\boldsymbol{\xi}}\|_\infty\}$.

Second, employing Lemma 3.5 in [Javanmard and Montanari \(2018\)](#), we know

$$\mathbb{P}(\phi_{\max}(k) \geq C_{\max} + C \sqrt{\frac{k}{n+N} + \frac{t}{\sqrt{n+N}}}) \leq 2 \exp(-ct^2 + k \log p + k), \quad (\text{S.23})$$

for $t \geq 0$, where C, c depend only on C_{\max} . Hence, $\|(\widehat{\boldsymbol{\Sigma}}_{n+N})_{\widehat{T}_2}\|_2 \leq \phi_{\max}(|\widehat{T}_2|) \leq \phi_{\max}(n+N) \leq c_1 \sqrt{\frac{(n+N) \log p}{n+N}}$ holds with large probability, and by inequality (S.22), we know $|\widehat{T}_2| \lesssim s\sqrt{\log p}$. Finally, we can refine the bound, as $\phi_{\max}(|\widehat{T}_2|) \leq \phi_{\max}(C's\sqrt{\log p}) \leq C_{\max}$ if $s(\log p)^2 = o(n+N)$. From (S.22), we derive that

$$|\widehat{T}_2| \leq \frac{64C_{\max}s}{(1-c_1)^2 C_{\min}}.$$

□

Now, we are ready to prove Theorem 4.1.

Proof. We take $\omega = 0$ in the basic inequality,

$$\frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T(\hat{\theta}_{SL} + \hat{\omega}))^2 + \lambda_\omega \|\hat{\omega}\|_1 \leq \frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T \hat{\theta}_{SL})^2.$$

Then we have

$$\begin{aligned} \frac{1}{2n} \sum_{i=1}^n (X_i^T \hat{\omega})^2 + \lambda_\omega \|\hat{\omega}\|_1 &\leq \frac{1}{n} \left| \sum_{i=1}^n (Y_i - X_i^T \hat{\theta}_{SL}) X_i^T \hat{\omega} \right| \\ &\leq \frac{1}{n} \left| \sum_{i=1}^n (Y_i - X_i^T \theta^*) X_i^T \hat{\omega} \right| + \frac{1}{n} \left| \sum_{i=1}^n X_i^T (\theta^* - \hat{\theta}_{SL}) X_i^T \hat{\omega} \right| \\ &\leq \left\| \frac{1}{n} \sum_{i=1}^n (Y_i - X_i^T \theta^*) X_i^T \right\|_\infty \|\hat{\omega}\|_1 + \frac{4 \sum_{i=1}^n (X_i^T (\theta^* - \hat{\theta}_{SL}))^2}{n} \\ &\quad + \frac{1}{4n} \sum_{i=1}^n (X_i^T \hat{\omega})^2. \end{aligned}$$

Denote $\lambda_0 = \left\| \frac{1}{n} \sum_{i=1}^n (Y_i - X_i^T \theta^*) X_i^T \right\|_\infty$. Choosing $\lambda_\omega \geq C_\omega K_1(\Phi + \sigma) \sqrt{\frac{\log p}{n}} \geq 2\lambda_0$, we then arrive that

$$\frac{1}{4n} \sum_{i=1}^n (X_i^T \hat{\omega})^2 + \frac{\lambda_\omega}{2} \|\hat{\omega}\|_1 \leq \frac{4}{n} \sum_{i=1}^n (X_i^T (\theta^* - \hat{\theta}_{SL}))^2.$$

With the given sparsity of $\text{supp}(\hat{\omega})$, we can proceed to show the rate of $\hat{\theta}_{SL}$ hold for $\hat{\theta}_R$ by

$$(1 - \delta) C_{\min} \|\hat{\omega}\|_2^2 \leq 16(1 + \delta) C_{\max} \|\theta^* - \hat{\theta}_{SL}\|_2^2.$$

That is to say $\|\hat{\theta}_R - \theta^*\|_2 \leq \|\theta^* - \hat{\theta}_{SL}\|_2 + \|\hat{\omega}\|_2 \leq (\frac{2C_{\max}}{C_{\min}} + 1) \|\theta^* - \hat{\theta}_{SL}\|_2$.

From the basic inequality, taking $\omega = \widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}$,

$$\frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T(\widehat{\boldsymbol{\theta}}_{SL} + \widehat{\boldsymbol{\omega}}))^2 + \lambda_\omega \|\widehat{\boldsymbol{\omega}}\|_1 \leq \frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L)^2 + \lambda_\omega \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}\|_1,$$

we can see

$$\frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_R)^2 \leq \frac{1}{2n} \sum_{i=1}^n (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L)^2 + \lambda_\omega (\|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}\|_1 - \|\widehat{\boldsymbol{\omega}}\|_1).$$

Therefore, we have

$$\frac{1}{2n} \sum_{i=1}^n (X_i^T(\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R))^2 \leq \frac{1}{n} \left| \sum_{i=1}^n (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L) X_i^T (\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R) \right| + \lambda_\omega \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_1.$$

We know with the same choice of $\lambda_\omega \geq CK_1(\Phi + \sigma) \sqrt{\frac{\log p}{n}}$ with large enough constant C ,

$$\lambda_\omega \geq \left\| \frac{\sum_{i=1}^n (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L) X_i}{n} \right\|_\infty = O_p(K_1(\Phi + \sigma) \sqrt{\frac{\log p}{n}}),$$

so we attain that

$$\frac{1}{2n} \sum_{i=1}^n (X_i^T(\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R))^2 \leq 2\lambda_\omega \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_1.$$

The support set of $\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R$ is contained in $\widehat{T}_1 \cup \widehat{T}_2$, so we have

$$\|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_1 \lesssim \sqrt{s} \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_2.$$

On the event of $\mathcal{B}_{c_1}(n, |\widehat{T}_1 \cup \widehat{T}_2|, 1)$,

$$\frac{1}{n} \sum_{i=1}^n (X_i^T(\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R))^2 \geq (1 - c_1)^2 C_{\min} \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_2^2.$$

Hence, $\|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_R\|_2 \leq \lambda_\omega \sqrt{s} / ((1 - c_1)^2 C_{\min})$. What is more, $\|\widehat{\boldsymbol{\theta}}_R - \boldsymbol{\theta}^*\|_2 \leq \|\widehat{\boldsymbol{\theta}}_R - \widehat{\boldsymbol{\theta}}_L\|_2 + \|\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*\|_2 \lesssim \lambda_\omega \sqrt{s}$.

□

S.3.5 Proof of Proposition 5.1

The proof is in a similar spirit as Theorem 4.1. We only provide a sketch.

Proof. Recall that the objective function can be rewritten as

$$(\hat{\boldsymbol{\theta}}_{h_1} + \boldsymbol{\omega})^T \hat{\boldsymbol{\Sigma}}_{n+N} (\hat{\boldsymbol{\theta}}_{h_1} + \boldsymbol{\omega}) - (\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)^T (\hat{\boldsymbol{\theta}}_{h_1} + \boldsymbol{\omega}) + 2\lambda_H \|\boldsymbol{\omega}\|_1.$$

where $\hat{\boldsymbol{\xi}}$ is defined with h_2 . First, we can take $\boldsymbol{\omega} = 0$ in the basic inequality, so

$$(\hat{\boldsymbol{\theta}}_{h_1} + \hat{\boldsymbol{\omega}})^T \hat{\boldsymbol{\Sigma}}_{n+N} (\hat{\boldsymbol{\theta}}_{h_1} + \hat{\boldsymbol{\omega}}) - (\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)^T \hat{\boldsymbol{\theta}}_R + 2\lambda_H \|\hat{\boldsymbol{\omega}}\|_1 \leq \hat{\boldsymbol{\theta}}_{h_1}^T \hat{\boldsymbol{\Sigma}}_{n+N} \hat{\boldsymbol{\theta}}_{h_1} - (\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)^T \hat{\boldsymbol{\theta}}_{h_1}.$$

Simplify it and we have

$$\begin{aligned} \hat{\boldsymbol{\omega}}^T \hat{\boldsymbol{\Sigma}}_{n+N} \hat{\boldsymbol{\omega}} + 2\lambda_H \|\hat{\boldsymbol{\omega}}\|_1 &\leq 2|\hat{\boldsymbol{\omega}}^T (\hat{\boldsymbol{\Sigma}}_{n+N} \hat{\boldsymbol{\theta}}_{h_1} - \frac{(\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)}{2})| \\ &\leq 2\|(\hat{\boldsymbol{\Sigma}}_{n+N} \boldsymbol{\theta}^* - \frac{(\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)}{2})\|_\infty \|\hat{\boldsymbol{\omega}}\|_1 + 2|\hat{\boldsymbol{\omega}}^T \hat{\boldsymbol{\Sigma}}_{n+N} (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1})| \\ &\leq \lambda_H \|\hat{\boldsymbol{\omega}}\|_1 + \frac{\hat{\boldsymbol{\omega}}^T \hat{\boldsymbol{\Sigma}}_{n+N} \hat{\boldsymbol{\omega}}}{2} + 2(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1})^T \hat{\boldsymbol{\Sigma}}_{n+N} (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1}) \end{aligned}$$

Hence we have

$$\frac{1}{2} \hat{\boldsymbol{\omega}}^T \hat{\boldsymbol{\Sigma}}_{n+N} \hat{\boldsymbol{\omega}} + \lambda_H \|\hat{\boldsymbol{\omega}}\|_1 \leq 2(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1})^T \hat{\boldsymbol{\Sigma}}_{n+N} (\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1})$$

As what we did previously, we choose $\lambda_H \geq 2\|\hat{\boldsymbol{\Sigma}}_{n+N} \boldsymbol{\theta}^* - \frac{(\hat{\boldsymbol{\xi}}_1 + \hat{\boldsymbol{\xi}}_2)}{2}\|_\infty$ as the one computing $\hat{\boldsymbol{\theta}}_{h_2}$.

Therefore, $\|\hat{\boldsymbol{\omega}}\|_2 \leq \sqrt{C_{\max}/C_{\min}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1}\|_2$ and $\|\hat{\boldsymbol{\theta}}_R - \boldsymbol{\theta}^*\|_2 \leq (\sqrt{C_{\max}/C_{\min}} + 1) \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{h_1}\|_2$.

On the other hand, we choose $\boldsymbol{\omega} = \widehat{\boldsymbol{\theta}}_{h_2} - \widehat{\boldsymbol{\theta}}_{h_1}$. Then

$$(\widehat{\boldsymbol{\theta}}_{h_1} + \widehat{\boldsymbol{\omega}})^T \widehat{\boldsymbol{\Sigma}}_{n+N} (\widehat{\boldsymbol{\theta}}_{h_1} + \widehat{\boldsymbol{\omega}}) - (\widehat{\boldsymbol{\xi}}_1 + \widehat{\boldsymbol{\xi}}_2)^T \widehat{\boldsymbol{\theta}}_R + 2\lambda_H \|\widehat{\boldsymbol{\omega}}\|_1 \leq \\ \widehat{\boldsymbol{\theta}}_{h_2}^T \widehat{\boldsymbol{\Sigma}}_{n+N} \widehat{\boldsymbol{\theta}}_{h_2} - (\widehat{\boldsymbol{\xi}}_1 + \widehat{\boldsymbol{\xi}}_2)^T \widehat{\boldsymbol{\theta}}_{h_2} + 2\lambda_H \|\widehat{\boldsymbol{\theta}}_{h_2} - \widehat{\boldsymbol{\theta}}_{h_1}\|_1.$$

It simplifies to

$$(\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2})^T \widehat{\boldsymbol{\Sigma}}_{n+N} (\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2}) + 2\lambda_H \|\widehat{\boldsymbol{\omega}}\|_1 \leq \\ 2|(\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2})^T (\widehat{\boldsymbol{\Sigma}}_{n+N} \widehat{\boldsymbol{\theta}}_{h_2} - \frac{(\widehat{\boldsymbol{\xi}}_1 + \widehat{\boldsymbol{\xi}}_2)}{2})| + 2\lambda_H \|\widehat{\boldsymbol{\theta}}_{h_2} - \widehat{\boldsymbol{\theta}}_{h_1}\|_1$$

Therefore, choose $\lambda_H \geq 2\|\widehat{\boldsymbol{\Sigma}}_{n+N} \widehat{\boldsymbol{\theta}}_{h_2} - \frac{(\widehat{\boldsymbol{\xi}}_1 + \widehat{\boldsymbol{\xi}}_2)}{2}\|_\infty$, and we have

$$(\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2})^T \widehat{\boldsymbol{\Sigma}}_{n+N} (\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2}) \leq 3\lambda_H \|\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2}\|_1 \quad (\text{S.24})$$

With the sparsity on \widehat{H}_1 and \widehat{H}_2 , we have $\|\widehat{\boldsymbol{\theta}}_{AH} - \widehat{\boldsymbol{\theta}}_{h_2}\|_2 \leq 3\lambda_{h_2}\sqrt{s}/((1-c_1)^2 C_{\min})$ and $\|\widehat{\boldsymbol{\theta}}_{AH} - \boldsymbol{\theta}^*\|_2 \leq 4\lambda_{h_2}\sqrt{s}/((1-c_1)^2 C_{\min})$. \square

S.4 Supplementary Technical Results

S.4.1 $\mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2 \asymp s$

Assume that the true conditional mean function $f(X)$ has an additive form $f(X) = \sum_{k \in S} f_k(X_k)$ with $\mathbb{E}(f_k(X_k)) = 0$, where X_k is the k -th component of X and S is a subset of $\{1, \dots, p\}$ with $|S| = s$. We further assume that all the covariates X_1, \dots, X_p are mutually independent.

We have

$$\begin{aligned}
\mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2 &= \mathbb{E} \left\{ \sum_{k \in S} (f_k(X_k) - X_k \theta_k^*) - \sum_{k \notin S} X_k \theta_k^* \right\}^2 \\
&= \mathbb{E} \left\{ \sum_{k \in S} (f_k(X_k) - X_k \theta_k^*) \right\}^2 + \mathbb{E} \left\{ \sum_{k \notin S} X_k \theta_k^* \right\}^2 \\
&= \sum_{k \in S} \mathbb{E}(f_k(X_k) - X_k \theta_k^*)^2 + \sum_{k \notin S} \Sigma_{kk}(\theta_k^*)^2,
\end{aligned}$$

where we use the fact that $\mathbb{E}(f_k(X_k)) = 0$ and $\mathbb{E}(X_k) = 0$. By the definition of $\boldsymbol{\theta}^*$, we know that for $k \notin S$, $\theta_k^* = 0$. For $k \in S$, $\theta_k^* = \operatorname{argmin}_{\theta} \mathbb{E}(f_k(X_k) - X_k \theta)^2 = \mathbb{E}(f_k(X_k)X_k)/\Sigma_{kk}$ and

$$\mathbb{E}(f_k(X_k) - X_k \theta_k^*)^2 = \mathbb{E}(f_k^2(X_k)) - [\mathbb{E}(f_k(X_k)X_k)]^2/\Sigma_{kk}$$

which is a constant. For example, if $X_k \sim N(0, 1)$ and $f_k(X_k) = X_k^3$, then $\mathbb{E}(f_k(X_k) - X_k \theta_k^*)^2 = 6$. Thus, we have $\mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2 = 6s$.

S.4.2 Rate of $\hat{\theta}_U$

Recall that the modified Dantzig selector $\hat{\theta}_U$ is defined as

$$\hat{\theta}_U = \arg \min \|\boldsymbol{\theta}\|_1, \quad \text{s.t.} \quad \|\hat{\Sigma}_{n+N} \boldsymbol{\theta} - \frac{1}{n} \sum_{i=1}^n X_i Y_i\|_{\infty} \leq \lambda_U.$$

Assume that Assumption 3.1 holds. By choosing $\lambda_U \asymp K_1 \sqrt{\frac{(\sigma^2 + \Phi^2 + \boldsymbol{\theta}^{*T} \boldsymbol{\Sigma} \boldsymbol{\theta}^*) \log p}{n}}$, we obtain that

$$\|\hat{\theta}_U - \boldsymbol{\theta}^*\|_1 = O_p \left(s K_1 (\Phi + \sigma + (\boldsymbol{\theta}^{*T} \boldsymbol{\Sigma} \boldsymbol{\theta}^*)^{1/2}) \sqrt{\frac{\log p}{n}} \right).$$

Under the conditions in our Theorem 3.2, we have $K_1 = O(1)$. This implies the bound (3.11).

The proof follows from the same argument as in Lemma S.4. The only nontrivial step is

to bound $\|\widehat{\Sigma}_{n+N}\boldsymbol{\theta} - \frac{1}{n} \sum_{i=1}^n X_i Y_i\|_\infty$. Using the triangle inequality, we have

$$\begin{aligned} \|\widehat{\Sigma}_{n+N}\boldsymbol{\theta} - \frac{1}{n} \sum_{i=1}^n X_i Y_i\|_\infty &\leq \left\| \frac{1}{n} \sum_{i=1}^n X_i (Y_i - f(X_i)) \right\|_\infty \\ &+ \left\| \frac{1}{n} \sum_{i=1}^n X_i (f(X_i) - X_i^T \boldsymbol{\theta}^*) \right\|_\infty + \|(\widehat{\Sigma}_n - \Sigma)\boldsymbol{\theta}^*\|_\infty + \|(\widehat{\Sigma}_{n+N} - \Sigma)\boldsymbol{\theta}^*\|_\infty. \end{aligned}$$

We have already derived in (S.17) that,

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n X_i (Y_i - f(X_i)) \right\|_\infty &\lesssim K_1 \sigma \sqrt{\frac{\log p}{n}}, \\ \left\| \frac{1}{n} \sum_{i=1}^n X_i (f(X_i) - X_i^T \boldsymbol{\theta}^*) \right\|_\infty &\lesssim K_1 \Phi \sqrt{\frac{\log p}{n}}. \end{aligned}$$

To control $\|(\widehat{\Sigma}_n - \Sigma)\boldsymbol{\theta}^*\|_\infty$, we note that $\|X_i\|_\infty \leq K_1$ and $\mathbb{E}(X_i^T \boldsymbol{\theta}^*)^2 = \boldsymbol{\theta}^{*T} \Sigma \boldsymbol{\theta}^*$. We obtain

$$\|(\widehat{\Sigma}_n - \Sigma)\boldsymbol{\theta}^*\|_\infty \lesssim K_1 (\boldsymbol{\theta}^{*T} \Sigma \boldsymbol{\theta}^*)^{1/2} \sqrt{\frac{\log p}{n}},$$

by the Nemirovski moment inequality in Lemma S.2 and Markov inequality. The last term $\|(\widehat{\Sigma}_{n+N} - \Sigma)\boldsymbol{\theta}^*\|_\infty$ is dominated by $\|(\widehat{\Sigma}_n - \Sigma)\boldsymbol{\theta}^*\|_\infty$ and can be ignored. Thus, we obtain

$$\|\widehat{\Sigma}_{n+N}\boldsymbol{\theta} - \frac{1}{n} \sum_{i=1}^n X_i Y_i\|_\infty \lesssim K_1 (\Phi + \sigma + (\boldsymbol{\theta}^{*T} \Sigma \boldsymbol{\theta}^*)^{1/2}) \sqrt{\frac{\log p}{n}}.$$

S.4.3 Rate of $\widehat{\theta}_T$

Although the method in Alquier and Hebiri (2012) does not apply cross-fitting procedure and they did not have a convergence rate result, to facilitate rate analysis here, we still use cross-fitting version of pseudo outcomes and let $\widetilde{\mathbf{Y}} = (\widehat{h}^{-j_1}(X_1), \dots, \widehat{h}^{-j_{n+N}}(X_{n+N}))$, where $j_i = 1, 2$ depending on which split of data X_i is in.

Recall that the transductive Dantzig selector is given by

$$\widehat{\boldsymbol{\theta}}_T = \operatorname{argmin} \|\boldsymbol{\theta}\|_1 \text{ s.t. } \frac{1}{n+N} \|\widetilde{\mathbf{X}}(\widetilde{\mathbf{Y}} - \widetilde{\mathbf{X}}^T \boldsymbol{\theta})\|_\infty \leq \lambda_T, \quad (\text{S.25})$$

where $\tilde{\mathbf{Y}} = (\tilde{Y}_1, \dots, \tilde{Y}_{n+N}) = (\hat{h}^{-j_1}(X_1), \dots, \hat{h}^{-j_{n+N}}(X_{n+N}))$. Under assumption 3.1 and conditions in Theorem 3.2, by choosing $\lambda_T \asymp K_1 \left(\Phi \sqrt{\frac{\log p}{n+N}} + (b_n + G_h) \right)$, we obtain that

$$\|\hat{\boldsymbol{\theta}}_T - \boldsymbol{\theta}^*\|_q = O_p \left(K_1 s^{1/q} \left\{ \Phi \sqrt{\frac{\log p}{n+N}} + (b_n + G_h) \right\} \right).$$

This rate is no better than our proposed methods when $b_n + G_h \gg (b_n + G_h + \sigma) \sqrt{\frac{\log p}{n}}$. This is true under the situation that $G_h = 0$ and b_n achieved with some traditional nonparametric rate such as the one discussed in Remark 3.3.

The proof mainly follows the same argument as in Lemma S.4. The only nontrivial step is to bound $\frac{1}{n+N} \|\tilde{\mathbf{X}}(\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}^T \boldsymbol{\theta})\|_\infty$. Using triangle inequality, we have

$$\frac{1}{n+N} \|\tilde{\mathbf{X}}(\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}^T \boldsymbol{\theta})\|_\infty \leq \left\| \frac{\sum_{i=1}^{n+N} X_i (\tilde{Y}_i - f(X_i))}{n+N} \right\|_\infty + \left\| \frac{\sum_{i=1}^{n+N} X_i (f(X_i) - X_i^T \boldsymbol{\theta}^*)}{n+N} \right\|_\infty$$

We have already derived in S.17 that,

$$\left\| \frac{1}{n+N} \sum_{i=1}^{n+N} X_i (f(X_i) - X_i^T \boldsymbol{\theta}^*) \right\|_\infty \lesssim K_1 \Phi \sqrt{\frac{\log p}{n+N}}.$$

To deal with $\left\| \frac{\sum_{i=1}^{n+N} X_i (\tilde{Y}_i - f(X_i))}{n+N} \right\|_\infty$, we adopt the other formulation of this term, which is

$$\begin{aligned} & \left\| \frac{\sum_{j=1}^2 \sum_{i \in D_j} X_i (\hat{h}^{-j}(X_i) - f(X_i))}{n+N} \right\|_\infty \leq \frac{1}{2} \sum_{j=1}^2 \left\| \frac{\sum_{i \in D_j} X_i (\hat{h}^{-j}(X_i) - f(X_i))}{n_j + N_j} \right\|_\infty \\ & \leq \frac{1}{2} \left(\sum_{j=1}^2 \left\| \frac{\sum_{i \in D_j} X_i (\hat{h}^{-j}(X_i) - f(X_i))}{n_j + N_j} - \mathbb{E}_{D_{-j}^*} [(\hat{h}^{-j}(X) - f(X))X] \right\|_\infty \right. \\ & \quad \left. + \left\| \mathbb{E}_{D_{-j}^*} [(\hat{h}^{-j}(X) - f(X))X] \right\|_\infty \right) \\ & = O_p \left(K_1 (b_n + G_h) \sqrt{\frac{\log p}{n+N}} + K_1 (b_n + G_h) \right) = O_p \left(K_1 (b_n + G_h) \right), \end{aligned}$$

where the bounds are derived similarly to S.12 and S.13. Thus, we obtain that

$$\frac{1}{n+N} \|\widetilde{\mathbf{X}}(\widetilde{\mathbf{Y}} - \widetilde{\mathbf{X}}^T \boldsymbol{\theta})\|_\infty \lesssim K_1 \left(\Phi \sqrt{\frac{\log p}{n+N}} + (b_n + G_h) \right).$$

S.4.4 Counter Example of Weighted Estimator

Since Nemirovski (2000), there is a vast literature on aggregation of estimators; see Rigollet et al. (2012) for a recent review. In the context of Section 4, as an alternative, we may consider the weighted estimator of the form $\widehat{\boldsymbol{\theta}}_\beta = \beta \widehat{\boldsymbol{\theta}}_{SL} + (1 - \beta) \widehat{\boldsymbol{\theta}}_L$ for some $\beta \in [0, 1]$, which corresponds to the convex combination of two estimators $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$. As suggested by Nemirovski (2000), to find the optimal weighted estimator, we have to split the sample. The first sample is used to compute the initial estimators $\widehat{\boldsymbol{\theta}}_{SL}$ and $\widehat{\boldsymbol{\theta}}_L$, while the optimal weight $\widehat{\beta}$ is estimated from the second sample, say $\{(X_1, Y_1), \dots, (X_{n'}, Y_{n'})\}$, by solving the following constrained least square problem

$$\widehat{\beta} = \underset{\beta \in [0, 1]}{\operatorname{argmin}} \frac{1}{n'} \sum_{i=1}^{n'} (Y_i - X_i^T (\beta \widehat{\boldsymbol{\theta}}_{SL} + (1 - \beta) \widehat{\boldsymbol{\theta}}_L))^2. \quad (\text{S.26})$$

The optimal weighted estimator is given by $\widehat{\boldsymbol{\theta}}_{\widehat{\beta}}$. However, our preliminary analysis suggests that in the semi-supervised setting, an optimal performance in the prediction does not necessarily lead to a more accurate estimation of $\boldsymbol{\theta}^*$. In other words, our Theorem 4.1 does not generally hold for the weighted estimator $\widehat{\boldsymbol{\theta}}_{\widehat{\beta}}$. To illustrate this argument, a counter example is given below. Along this line, one future research question is that, in the context of Section 5, whether it is possible to combine an increasing number of semi-supervised estimators via some more refined aggregation methods, such as Q-aggregation (Dai et al., 2012).

We next study the property of this weighted estimator. Above is a univariate least square problem and denote

$$\widehat{W} = \frac{\sum_{i=1}^{n'} (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L) X_i^T (\widehat{\boldsymbol{\theta}}_{SL} - \widehat{\boldsymbol{\theta}}_L)}{\sum_{i=1}^{n'} (X_i^T (\widehat{\boldsymbol{\theta}}_{SL} - \widehat{\boldsymbol{\theta}}_L))^2}. \quad (\text{S.27})$$

Then

$$\widehat{\beta} = \begin{cases} \widehat{W} & \text{if } \widehat{W} \in [0, 1], \\ 0 & \text{if } \widehat{W} < 0, \\ 1 & \text{if } \widehat{W} > 1. \end{cases}$$

With $\widehat{\beta}$, we denote $\widehat{\boldsymbol{\theta}}_{\widehat{\beta}} = \widehat{\beta} \widehat{\boldsymbol{\theta}}_{SL} + (1 - \widehat{\beta}) \widehat{\boldsymbol{\theta}}_L$.

Firstly, we consider the case where $\widehat{\beta} = 0$, and $\widehat{\beta} = 1$ follows the same argument. When $\widehat{\beta} = 0$, we have

$$\sum_{i=1}^{n'} (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_L)^2 \leq \sum_{i=1}^{n'} (Y_i - X_i^T \widehat{\boldsymbol{\theta}}_{SL})^2, \quad (\text{S.28})$$

from which we can get

$$(\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*)^T \widehat{\Sigma}_n (\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*) \leq (\widehat{\boldsymbol{\theta}}_{SL} - \boldsymbol{\theta}^*)^T \widehat{\Sigma}_n (\widehat{\boldsymbol{\theta}}_{SL} - \boldsymbol{\theta}^*) + \frac{2}{n'} \sum_{i=1}^{n'} ((\eta_i + \epsilon_i) X_i^T (\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL})). \quad (\text{S.29})$$

The remainder term $\frac{2}{n'} \sum_{i=1}^{n'} ((\eta_i + \epsilon_i) X_i^T (\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}))$ is $O_p(\frac{K_1(\Phi+\sigma) \|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}\|_2}{\sqrt{n}})$. Intuitively, by triangle inequality, we can show that $\|\widehat{\boldsymbol{\theta}}_L - \widehat{\boldsymbol{\theta}}_{SL}\|_2$ is of order $K_1(\sigma + \Phi) \sqrt{\frac{s \log p}{n}}$, and therefore the remainder term is $O_p(\frac{K_1^2(\Phi+\sigma)^2 \sqrt{s \log p}}{n})$. However, the inequality S.29 shows that $\|\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*\|_2$ is not guaranteed to be of smaller order than $\|\widehat{\boldsymbol{\theta}}_{SL} - \boldsymbol{\theta}^*\|_2$, although the prediction from $\widehat{\boldsymbol{\theta}}_L$ is preferred.

Indeed, in this case $\|\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*\|_2$ can be larger than the order of $\|\widehat{\boldsymbol{\theta}}_{SL} - \boldsymbol{\theta}^*\|_2$. To illustrate this argument, a counter example is given as follows.

Assume that the true conditional mean function $f(X)$ is supported on S , a subset of $\{1, \dots, p\}$ with $|S| = s$ and projection coefficients are nonzero only on S with value θ . We further assume that all the covariates X_1, \dots, X_p are i.i.d. uniformly distributed on $[-1, 1]$.

Suppose we have two initial estimators $\widehat{\boldsymbol{\theta}}_L$ and $\widehat{\boldsymbol{\theta}}_{SL}$ whose support sets are both $S \cup S'$, where $|S'| = s$. Let $\widehat{\boldsymbol{\theta}}_L^k = \widehat{\boldsymbol{\theta}}_{SL}^k = \theta + \delta$ for $k \in S$, $\widehat{\boldsymbol{\theta}}_L^k = \delta_1$, $\widehat{\boldsymbol{\theta}}_{SL}^k = \delta_2$ for $k \in S'$ and $\delta_1, \delta_2 \geq 0$. Next, we construct examples where $\widehat{\boldsymbol{\theta}}_L$ achieves the minimal prediction

error among all the convex weighted estimator but has arbitrarily larger order of estimation error than $\widehat{\boldsymbol{\theta}}_{SL}$. We can compute \widehat{W} according to equation (S.27), and the denominator is $\sum_{i=1}^{n'} (\sum_{k \in S'} X_i^k)^2 (\delta_1 - \delta_2)^2$, where X_i^k is the k -th coordinate of X_i . And the numerator is $\sum_{i=1}^{n'} (Y_i - \sum_{k \in S} X_i^k (\theta - \delta) - \sum_{k \in S'} X_i^k \delta_1) (\sum_{k \in S'} X_i^k (\delta_2 - \delta_1))$. Set $\delta = \sigma \sqrt{\frac{\log p}{n'}}$, $\delta_2 = \Phi \sqrt{\frac{\log p}{n' + N}}$ and $\delta_1 = \rho \sqrt{\frac{\log p}{n'}}$, where $N \gg n'$ and $\rho \ll \Phi$ but $\delta_1 \gg \delta_2$, so that $\|\widehat{\boldsymbol{\theta}}_L - \boldsymbol{\theta}^*\|_2 \gg \|\widehat{\boldsymbol{\theta}}_{SL} - \boldsymbol{\theta}^*\|_2^2$.

Denote

$$T_1 = \frac{1}{n'} \sum_{i=1}^{n'} (\sum_{k \in S'} X_i^k)^2 (\delta_1 - \delta_2)^2,$$

$$T_2 = \frac{1}{n'} \sum_{i=1}^{n'} (Y_i - \sum_{k \in S} X_i^k (\theta - \delta)) (\sum_{k \in S'} X_i^k (\delta_2 - \delta_1)),$$

and

$$T_3 = \frac{1}{n'} \sum_{i=1}^{n'} (\sum_{k \in S'} X_i^k \delta_1) (\sum_{k \in S'} X_i^k (\delta_1 - \delta_2)),$$

so that the denominator is T_1 and the numerator is $T_2 - T_3$. We firstly notice that $\mathbb{E}[T_1] = \frac{1}{3} s (\delta_1 - \delta_2)^2$, and $\mathbb{E}[T_3] = \frac{1}{3} s (\delta_1 - \delta_2) \delta_1$. Although $\mathbb{E}[T_2] = 0$, but we will show that T_2 can be as large as the order $s \delta_1 (\delta_1 - \delta_2)$. By Chebyshev's inequality,

$$\Pr \left(\left| \frac{1}{n'} \sum_{i=1}^{n'} (Y_i - \theta \sum_{k \in S} X_i^k) (\sum_{k \in S'} X_i^k (\delta_2 - \delta_1)) \right| > t \right) \leq \frac{\mathbb{E}[(Y - \theta \sum_{k \in S} X^k)^2] \mathbb{E}[(\sum_{k \in S'} X^k)^2] (\delta_2 - \delta_1)^2}{n' t^2}$$

$$\leq \frac{s (\delta_2 - \delta_1)^2 (\Phi^2 + \sigma^2)}{n' t^2}, \quad (\text{S.30})$$

so we have $\frac{1}{n'} \sum_{i=1}^{n'} (Y_i - \theta \sum_{k \in S} X_i^k) (\sum_{k \in S'} X_i^k (\delta_2 - \delta_1)) = O_p(\frac{\sqrt{s(\Phi^2 + \sigma^2)} (\delta_1 - \delta_2)}{\sqrt{n'}})$. Therefore, we can see that as long as $\sqrt{\frac{\Phi^2 + \sigma^2}{n'}} \gg \sqrt{s} \delta_1 = \rho \sqrt{\frac{s \log p}{n'}}$, or with another sufficient condition that $\Phi^2 \gg s \rho^2 \log p$, $|T_2|$ can be as large as T_3 . Then there is a non-vanishing probability that $T_2 + T_3 \leq 0$. In this case, the optimal weight $\widehat{\beta}$ is 0, which indicates $\widehat{\boldsymbol{\theta}}_L$ is more favorable. However, it is clear to see that the estimation error of $\widehat{\boldsymbol{\theta}}_L$ is of larger order than that of $\widehat{\boldsymbol{\theta}}_{SL}$.

At last, we show an example where $\Phi^2 \gg s \rho^2 \log p$. Set $f(X) = (\sum_{k \in S} X^k)^3$, and denote $A = \frac{\sum_{k \in S} X^k}{\sqrt{s}}$, then we know $\text{Var}(A) = \frac{1}{3}$ and $\sqrt{s} A$ follows a Irwin-Hall Distribution so that $\sqrt{3} A$ can be approximate well by a standard Gaussian variable. Hence,

$\theta = \mathbb{E}[(\sum_{k \in S} X^k)^4]/\mathbb{E}[(\sum_{k \in S} X^k)^2] = s \frac{\mathbb{E}[A^4]}{\mathbb{E}[A^2]}$ should be of order s and $\Phi^2 = \mathbb{E}[(\sum_{k \in S} X^k)^3 - \theta(\sum_{k \in S} X^k)^2] = s^3 \mathbb{E}[(A^3 - \frac{\theta}{s} A^2)]$ is of order s^3 .

S.4.5 Further intuition on modified score function

Recall that the score function of the supervised lasso/Dantzig selector can be rewritten as

$$S_1 = \frac{1}{n} \sum_{i=1}^n X_i(Y_i - X_i^T \boldsymbol{\theta}^*) = \frac{1}{n} \sum_{i=1}^n X_i(Y_i - f(X_i)) + \frac{1}{n} \sum_{i=1}^n X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*).$$

With unlabeled data, we propose to replace the last term with $\frac{1}{n+N} \sum_{i=1}^{n+N} X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*)$, the sample average over both labeled and unlabeled data. This leads to the following modified score function

$$S_2 = \frac{1}{n} \sum_{i=1}^n X_i(Y_i - f(X_i)) + \frac{1}{n+N} \sum_{i=1}^{n+N} X_i(f(X_i) - X_i^T \boldsymbol{\theta}^*). \quad (\text{S.31})$$

To see why this manipulation is useful without going through the technical details (e.g, concentration inequalities), let us compare the variance of S_1 and S_2 . Under mild assumptions (e.g. data are i.i.d, ϵ and X are independent), it is easily seen that

$$\text{Cov}(S_1) = \frac{1}{n}(\sigma^2 + \Phi^2)\boldsymbol{\Sigma}, \text{ and } \text{Cov}(S_2) = \frac{1}{n}\sigma^2\boldsymbol{\Sigma} + \frac{1}{n+N}\Phi^2\boldsymbol{\Sigma},$$

where $\sigma^2 = \text{Var}(\epsilon)$, $\Phi^2 = \mathbb{E}(f(X) - X^T \boldsymbol{\theta}^*)^2$ and $\boldsymbol{\Sigma} = \text{Cov}(X)$. Clearly, $\text{Cov}(S_1) - \text{Cov}(S_2)$ is always a semi-positive definite matrix, which implies that the variability of the proposed modified score function is smaller than the original score function. In particular, by comparing the two terms $\frac{1}{n+N}\Phi^2\boldsymbol{\Sigma}$ and $\frac{1}{n}\Phi^2\boldsymbol{\Sigma}$ in $\text{Cov}(S_2)$ and $\text{Cov}(S_1)$ respectively, we can claim that our manipulation can reduce the variance due to model misspecification.

S.5 Examples of \hat{h}

Sparse additive models. Assume that Y_i given $X_i = (X_{i1}, \dots, X_{ip})$ follows the additive

model

$$Y_i = \mu + \sum_{j=1}^p f_j(X_{ij}) + \epsilon_i,$$

where μ is an intercept, f_j 's are unknown functions and ϵ_i is the random error. For identification purpose, we assume $\mathbb{E}f_j(X_{ij}) = 0$. Suppose that some of the functions f_j are 0. We also assume f_j is $\bar{\ell}$ -smooth, that is the k th derivative $f_j^{(k)}$ exists and satisfies the Lipschitz condition $|f_j^{(k)}(t_1) - f_j^{(k)}(t_2)| \leq C|t_1 - t_2|^{\bar{\ell}-k}$, where C is a constant and k is the greatest integer strictly less than $\bar{\ell}$.

The literature on how to estimate these additive functions is vast, see [Lin and Zhang \(2006\)](#); [Meier et al. \(2009\)](#); [Huang et al. \(2010\)](#); [Raskutti et al. \(2012\)](#) among many others. Here, we adopt the adaptive group lasso approach in [Huang et al. \(2010\)](#) as a concrete example. Suppose X_j takes values in $[a, b]$. Let $a = c_0 < c_1 < \dots < c_{K+1} = b$ be a partition of $[a, b]$ into K subintervals $I_t = [c_t, c_{t+1})$, where $K = n^v$ for some $0 < v < 1/2$ and $\max_k |c_k - c_{k-1}| = O(n^{-v})$. Let S_n denote the space of polynomial splines of order ℓ , that is the restriction of the function to I_t is a polynomial of degree ℓ and for $\ell \geq 2$ the function is $\ell - 2$ times continuously differentiable on $[a, b]$. For the space of splines S_n , there exists a normalized B-spline basis $\{\phi_k, 1 \leq k \leq m_n\}$, where $m_n = K + \ell$, such that for any $f \in S_n$ we have $f(x) = \sum_{k=1}^{m_n} \beta_k \phi_k(x)$.

Based on the spline theory, any $\bar{\ell}$ -smooth function f_j can be approximated by $\sum_{k=1}^{m_n} \beta_{jk} \phi_k(x)$. Thus, [Huang et al. \(2010\)](#) proposed to estimate the unknown functions by solving the following group lasso problem

$$\min \sum_{i=1}^n \left[Y_i - \mu - \sum_{j=1}^p \sum_{k=1}^{m_n} \beta_{jk} \phi_k(X_{ij}) \right]^2 + \lambda_n \sum_{j=1}^p w_j \|\boldsymbol{\beta}_j\|_2,$$

where λ_n is the tuning parameter, w_j is weight in the adaptive lasso, and $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jm_n})$. Let us denote the resulting estimator by $\hat{\mu}$ and $\hat{\beta}_{jk}$ for $1 \leq j \leq p$ and $1 \leq k \leq m_n$. Then we

have

$$\widehat{h}(x) = \widehat{\mu} + \sum_{j=1}^p \sum_{k=1}^{m_n} \widehat{\beta}_{jk} \phi_k(x_j).$$

Under the conditions detailed in [Huang et al. \(2010\)](#), their Corollary 2 implies that $\|\widehat{h} - f\|_2 = O_p(s^{1/2}n^{-\bar{\ell}/(2\bar{\ell}+1)})$, where $\bar{\ell}$ is the smoothness of the function f_j and s is the number of nonzero f_j 's.

Pairwise interaction models. [Zhao et al. \(2016\)](#) considered the following interaction models,

$$Y_i = \sum_{j=1}^p \gamma_j X_{ij} + \sum_{1 \leq j \leq k \leq p} \gamma_{jk} X_{ij} X_{ik} + \epsilon_i$$

where ϵ_i is the random error and $\{\gamma_j\}_{1 \leq j \leq p}$ and $\{\gamma_{jk}\}_{1 \leq j \leq k \leq p}$ are unknown parameters. Assume that $\{\gamma_j\}_{1 \leq j \leq p}$ and $\{\gamma_{jk}\}_{1 \leq j \leq k \leq p}$ are sparse. [Zhao et al. \(2016\)](#) proposed the following lasso estimator

$$\min \sum_{i=1}^n \left[Y_i - \sum_{j=1}^p \gamma_j X_{ij} + \sum_{1 \leq j \leq k \leq p} \gamma_{jk} X_{ij} X_{ik} \right]^2 + \lambda_1 \sum_{j=1}^p |\gamma_j| + \lambda_2 \sum_{1 \leq j \leq k \leq p} |\gamma_{jk}|.$$

Let us denote the resulting estimator by $\widehat{\gamma}_j$ and $\widehat{\gamma}_{jk}$. Then we have

$$\widehat{h}(x) = \sum_{j=1}^p \widehat{\gamma}_j x_j - \sum_{1 \leq j \leq k \leq p} \widehat{\gamma}_{jk} x_j x_k.$$

In this case, $\|\widehat{h} - f\|_2$ corresponds to the prediction error in the lasso problem. Following the proof in [Zhao et al. \(2016\)](#), it can be shown that $\|\widehat{h} - f\|_2 = O_p(\sqrt{s \log p/n})$, where s is the number of nonzero parameters.

We also note that in the analysis of interaction models, one common assumption is known as the strong hierarchy principle, that is if an interaction $X_j X_k$ exists then their main effects (the linear terms of X_j and X_k) also exist. [Zhao et al. \(2016\)](#) also proposed a group lasso approach to accommodate the strong hierarchy principle in interaction models.

	N=300		N=900		N=1500	
	L_1 error	L_2 error	L_1 error	L_2 error	L_1 error	L_2 error
SSL12	3.93	1.17	2.50	0.86	2.15	0.78
SSL21	3.84	1.19	2.40	0.86	2.10	0.79

Table 1: The L_1 and L_2 estimation errors of SSL12, SSL21 with Model 1 for $n = 300$, $p = 1000$.

S.6 Supplementary Simulation Results

Parallel to the results presented in the main paper for Models 1, 2, 3 with $p = 500$ and $n = 200$, here we also present the similar results for Models 1, 2, 3 with $p = 200$ and $n = 100$ in Figures 2, 3 and 4, as well as for Models 1, 2, 3 with $p = 1000$ and $n = 300$ in Figures 5, 6 and 7, respectively. The results from analyzing the L_2 and L_1 estimation errors are similar to the main paper, so are omitted.

When creating the aggregated estimator $\widehat{\boldsymbol{\theta}}_{AH}$ in Section 5, suppose that we have two candidate models h_1 and h_2 for the conditional mean functions. The aggregated estimator in (5.1) is given by $\widehat{\boldsymbol{\theta}}_{AH} = \widehat{\boldsymbol{\theta}}_{h_1} + \widehat{\boldsymbol{\omega}}$, where we use the semi-supervised estimator with model h_1 , $\widehat{\boldsymbol{\theta}}_{h_1}$, as an initial estimator. The creation of $\widehat{\boldsymbol{\theta}}_{AH}$ is not symmetric to h_1 and h_2 . Theoretically, Proposition 5.1 shows that the aggregated estimators have the same convergence rate regardless whether we use $\widehat{\boldsymbol{\theta}}_{h_1}$ or $\widehat{\boldsymbol{\theta}}_{h_2}$ as the initial estimator. To back up our theory, we add new simulation results for the aggregated estimator named as SSL21, which uses $\widehat{\boldsymbol{\theta}}_{h_2}$ as the initial estimator. Tables 1,2,3 show that the two aggregated estimators SSL21 and SSL12 have similar performance.

Additionally, in Section 3.2, we adopt the splitting strategy 50-50 in the cross-fitting technique; i.e., we use 50% of labeled data to train the model \widehat{h}^{-j} and the other 50% to compute $\widehat{\boldsymbol{\xi}}_j$; see (3.4). It is of interest to consider an imbalanced split, such as 70-30. Note that, due to the symmetry of calculating the estimator $\widehat{\boldsymbol{\theta}}_{SD}$ in (3.5), the split strategies of 70-30 and 30-70 essentially result in the same estimator. The comparison results in Table 4 show that the splitting strategy of 50-50 is slightly preferred.

Estimation Results for Model 1 ($p=200, n=100$)

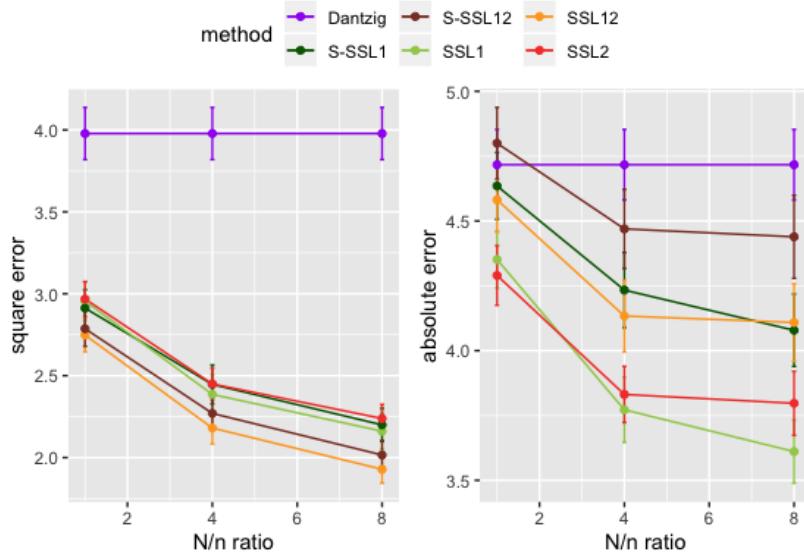


Figure 2: The L_2 and L_1 estimation error under Model 1 with $p = 200$ and $n = 100$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 9.14(0.32), 11.16(0.34), 11.47(0.35). L_1 errors for U-Dantzig are 7.68(0.20), 8.33(0.19), 8.49(0.19).

Estimation Results for Model 2 ($p=200, n=100$)

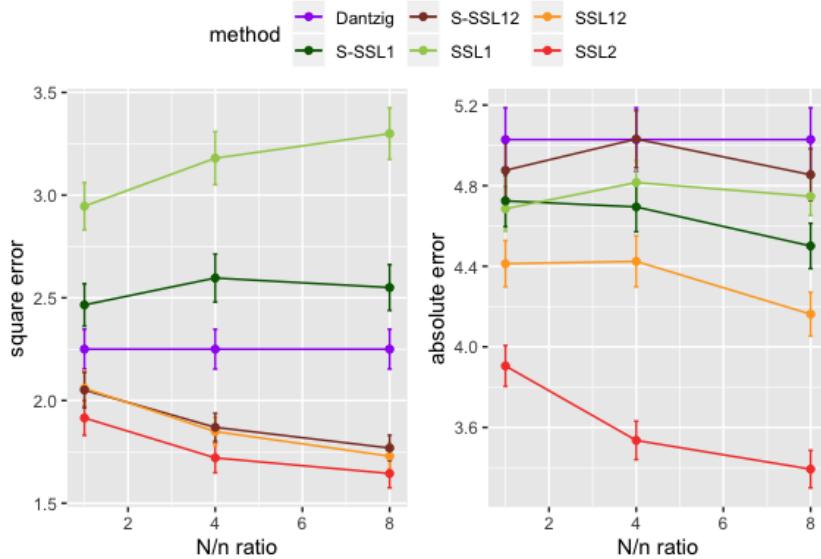


Figure 3: The L_2 and L_1 estimation error under Model 2 with $p = 200$ and $n = 100$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 4.21 (0.14), 4.25(0.14), 4.24(0.14). L_1 errors for U-Dantzig are 5.32(0.16), 5.29(0.15), 5.50(0.15).

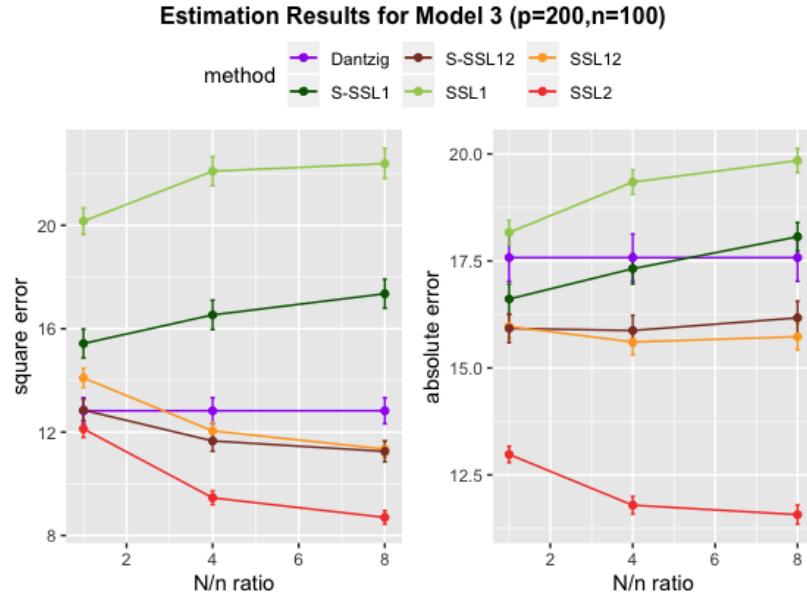


Figure 4: The L_2 and L_1 estimation error under Model 3 with $p = 200$ and $n = 100$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 17.46(0.45), 19.03(0.40), 19.73(0.44). L_1 errors for U-Dantzig are 16.99(0.27), 17.70(0.26), 17.88(0.31).

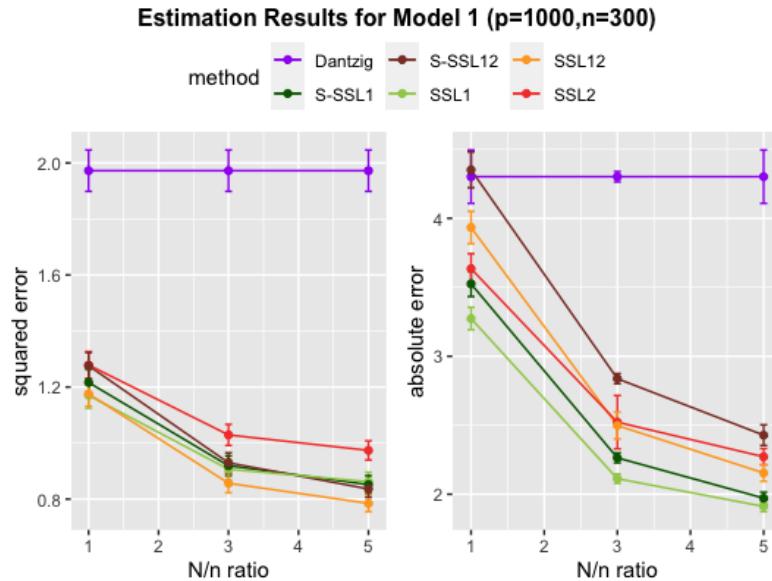


Figure 5: The L_2 and L_1 estimation error under Model 1 with $p = 1000$ and $n = 300$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are 4.12(0.12), 4.94(0.14), 5.18(0.15). L_1 errors for U-Dantzig are 5.09(0.11), 5.57(0.10), 5.74(0.12).

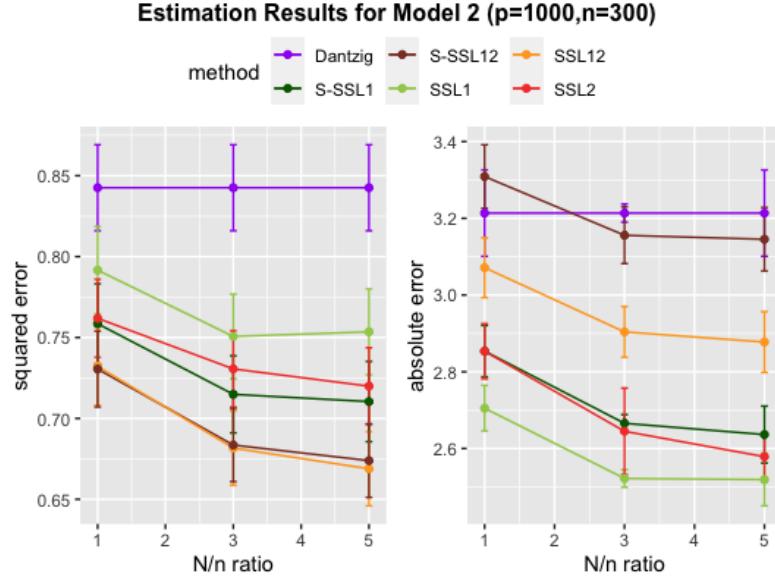


Figure 6: The L_2 and L_1 estimation error under Model 2 with $p = 1000$ and $n = 300$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are $1.82(0.06)$, $2.01(0.06)$, $2.11(0.07)$. L_1 errors for U-Dantzig are $3.47(0.06)$, $3.89(0.07)$, $3.98(0.07)$.

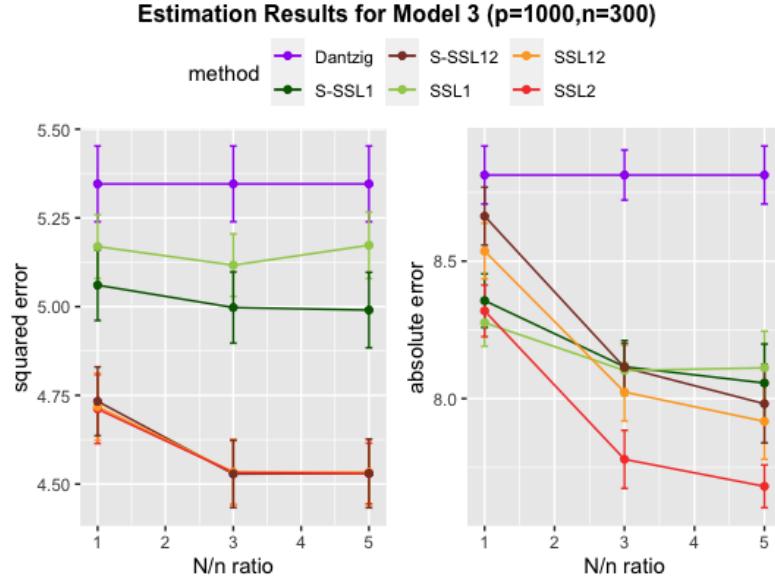


Figure 7: The L_2 and L_1 estimation error under Model 3 with $p = 1000$ and $n = 300$. The length of the vertical bar represents the magnitude of the sample standard deviations. L_2 errors for U-Dantzig are $8.82(0.18)$, $10.32(0.21)$, $10.77(0.21)$. L_1 errors for U-Dantzig are $11.81(0.14)$, $13.97(0.19)$, $14.33(0.20)$.

	N=300		N=900		N=1500	
	L_1 error	L_2 error	L_1 error	L_2 error	L_1 error	L_2 error
SSL12	3.07	0.73	2.90	0.68	2.88	0.67
SSL21	2.98	0.73	2.70	0.67	2.63	0.66

Table 2: The L_1 and L_2 estimation errors of SSL12, SSL21 with Model 2 for $n = 300$, $p = 1000$.

	N=300		N=900		N=1500	
	L_1 error	L_2 error	L_1 error	L_2 error	L_1 error	L_2 error
SSL12	8.54	4.72	8.02	4.53	7.92	4.53
SSL21	8.39	4.71	7.81	4.51	7.70	4.51

Table 3: The L_1 and L_2 estimation errors of SSL12, SSL21 with Model 3 for $n = 300$, $p = 1000$.

	split ratio	Model 1		Model 2		Model 3	
		SSL1	SSL2	SSL1	SSL2	SSL1	SSL2
L_2 error	50-50	0.74	1.03	0.80	0.65	5.65	4.25
	70-30	0.78	1.08	0.96	0.74	6.06	4.55
L_1 error	50-50	2.20	2.53	2.97	2.62	9.06	7.77
	70-30	2.44	2.74	3.55	2.62	9.30	8.07

Table 4: The L_1 and L_2 estimation errors of estimators SSL1 and SSL2 with split ratio 50-50 and 70-30, for the setting with $p = 1000$, $n = 400$ and $N = 1200$.

S.7 Supplementary Materials in Real Data Application

The MIMIC-III database is comprehensive in nature and it includes 26 individual tables in csv format, named “admissions”, “patients”, “outpuvents”, “chartevents”, etc. It includes vital signs, medications, laboratory measurements, observations and notes charted by care providers. The total data size reached 40GB after decompression. There are two types of data in the database: static data and dynamic data. The static data do not change over time and are recorded only once such as the date of birth, while the dynamic data are mainly the vital signs that periodically measured during the patient’s ICU stay such as the blood pressure and the heart rate. The Table 3 of [Johnson et al. \(2016\)](#) contains all the classes of the data available in MIMIC-III while their Table 4 provides an overview of the 26 data tables.

Due to the complexity of this large database, a few preprocessing steps are conducted to improve the quality of the data.

For the patient selection process, we exclude patients with less than 24-hour ICU stays, since those patients with short length of stays have many missing values in the covariates. Patients who died during the ICU stay are also excluded. Additionally, for patients with multiple ICU admissions, we only keep the first ICU admission in our analysis. This database has a variety of missing values with different patterns, including for the outcome “albumin”. To illustrate our method, we choose not to impute this outcome and instead only focus on the patients whose “albumin” is available.

We consult the previous literature ([Brown et al., 2012](#); [Tabak et al., 2017](#); [Xue et al., 2019](#); [Liu et al., 2020](#); [Du et al., 2021](#); [Jhou et al., 2021](#); [Tang et al., 2021](#)) to select the variables used in our data analysis. In general, the variable selection takes into account the following directives ([Fialho et al., 2012](#); [Nates et al., 2016](#)). First, variables must be easily and/or routinely assessed in the 24 hours before discharge. There is a special concern for the patients who were older than 89 years, because their dates of birth were shifted to obscure their true age to comply with Health Insurance Portability and Accountability Act

(HIPAA) regulations. These patients appear in the database with ages of over 300 years, and the MIMIC-III team only provides the median age for the patients before shifting, which is 91.4. Thus, for the patients with shifted age, we replace their age with the median 91.4. Second, the balance should be kept in the number of selected variables given that it will affect the number of patients that will be used in the analysis, i.e., the more variables that are defined, the fewer patients are likely to have all of them collected at the same time; or the high number of variables will bias the dataset towards selecting patients having similar conditions. Third, the variables selected should be independent with minimum correlation. For dynamic data such as many clinical biomarkers with continuous scale, the database collects the minimum, the maximum, as well as the mean, values during each patient's ICU stay. In order to alleviate the potential collinearity among these variables, we decide to only include the mean values in our analysis. Additionally, besides the 78 covariates themselves, we include pairwise interactions among 75 of the covariates that are continuous-scale clinical biomarkers as well as their squared terms, which all together result totally $p = 2928$ variables.

Regarding other missing data issues, in our situation, around 54% covariates contain missing values. Among these covariates with missing values, the missingness proportions are 9.4% on average and the range is from 0.2% to 30.8%. For the demographic variables, most observations were complete. For the laboratory test observations, observations are missing mainly because the variables are considered to be irrelevant to the current clinical problems, thus those laboratory tests are not requested by clinician. For those missing values, we simply impute them using the mean of observed samples, the so-called mean imputation.

After all of these data pre-processing steps, we are left with a dataset with 4784 patients and each of them has $p = 2928$ covariates. In our analysis, we consider a high dimensional setting where the sample size of our labeled data is $n = 500$.

Besides the results reported in the main text, we also carry out a parallel analysis and summarize the results, aggregating from 100 replications, in Figure 8. The difference is, here we set aside 2500 patients that are not going to be used in the computation of the three

Supplement: Estimation Results for Real Data

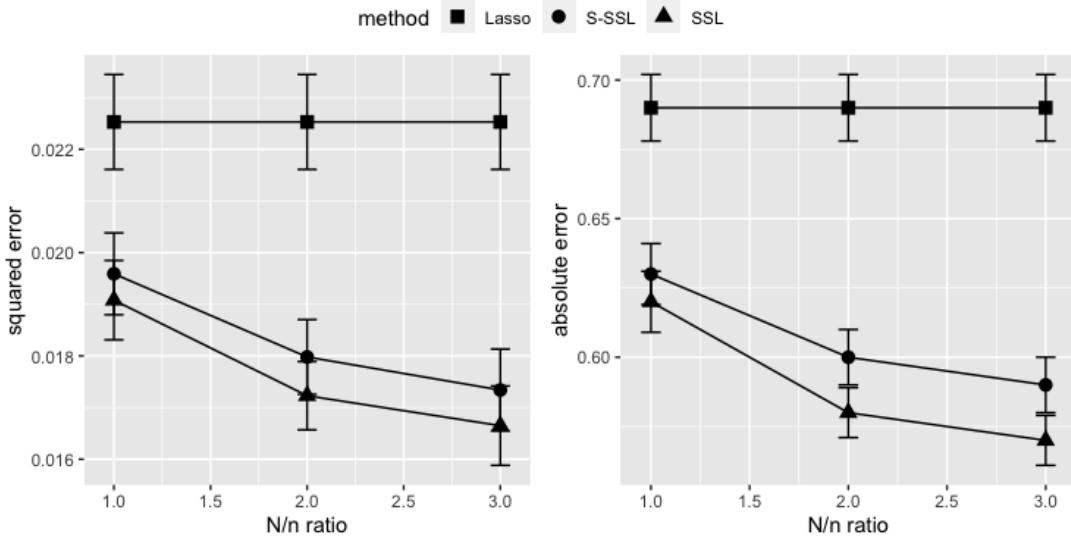


Figure 8: The L_2 and L_1 estimation error for real data application reported in the supplement.

comparative estimators. Instead, these 2500 patients are used to compute the lasso estimator, regarded as the proxy of the underlying true linear coefficients of the working model in this situation. In each replication we randomly choose 2000 patients from the remaining 2284 patients and call the first $n = 500$ patients the labeled data. The corresponding results in Figure 8 show a similar pattern as in the main text.

Last but not least, regarding our analysis results, we also collect the variable selection results from the comparison of supervised lasso, S-SSL and SSL estimators. The results show that the variables `white blood cells`, `hemoglobin` and `lymphocytes` are selected more frequently by both of SSL and S-SSL than supervised lasso. In the medical literature, there are various evidence documenting the associations between the albumin level in the blood sample with our selected biomarkers here, such as with the hemoglobin level (Fukui et al., 2008), with the lymphocytes level (Alagappan et al., 2018), and with the count of the white blood cells (Cavalot et al., 2002). More interestingly, researchers recently started to use the so-called HALP score, a combination of hemoglobin, albumin, lymphocyte, and platelet levels in the blood sample (Chen et al., 2015), as a prognostic factor for various disease types

(Peng et al., 2018; Guo et al., 2019; Shen et al., 2019; Tojek et al., 2019). By exploring the data with the random forest model, our statistical results from the semi-supervised methods are consistent with the medical literature.

S.8 Discussion

In this paper, we consider high-dimensional semi-supervised learning, and focus on when and how we can exploit the unlabeled data to achieve the optimal and safe estimates of model parameters. Our key observation is that, the covariate X in the unlabeled data, whose sample size can be much bigger than the labeled data, could be very informative for the parameter of interest in a misspecified model. We derive the minimax lower bound for parameter estimation in the semi-supervised setting, and show that generally the traditional supervised estimators without using the unlabeled data cannot attain this lower bound. Then, we propose a semi-supervised estimator, which depends on the correct specification of the conditional mean function, can attain the minimax lower bound and hence is optimal. To alleviate the strong requirement for this optimal estimator, we further propose a safe semi-supervised estimator. We view it safe, because this estimator remains minimax optimal when the conditional mean function is correctly specified, and is always at least as good as the supervised estimators.

While our focus is on the parameter estimation, we should note some recent and important progress in inference, particularly in the construction of confidence intervals with high-dimensional data in the semi-supervised setting (Zhang and Bradic, 2019; Cai and Guo, 2020). It would be important and challenging to develop procedures for statistical inference under our framework.

References

Alagappan, M., Pollom, E. L., von Eyben, R., Kozak, M. M., Aggarwal, S., Poulsides, G. A., Koong, A. C., and Chang, D. T. (2018), “Albumin and neutrophil-lymphocyte ratio (NLR) predict survival in patients with pancreatic adenocarcinoma treated with SBRT,” *American journal of clinical oncology*, 41, 242–247.

Alquier, P. and Hebiri, M. (2012), “Transductive versions of the LASSO and the Dantzig Selector,” *Journal of Statistical Planning and Inference*, 142, 2485–2500.

Bhlmann, P. and Van de Geer, S. (2011), *Statistics for High-Dimensional Data: Methods, Theory and Applications*, Springer Publishing Company, Incorporated, 1st ed.

Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009), “Simultaneous analysis of Lasso and Dantzig selector,” *Ann. Statist.*, 37, 1705–1732.

Brown, S. E., Ratcliffe, S. J., Kahn, J. M., and Halpern, S. D. (2012), “The epidemiology of intensive care unit readmissions in the United States,” *American Journal of Respiratory and Critical Care Medicine*, 185, 955–964.

Cai, T. and Guo, Z. (2020), “Semisupervised inference for explained variance in high dimensional linear regression and its applications,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82, 391–419.

Cavalot, F., Massucco, P., Perna, P., Traversa, M., Anfossi, G., and Trovati, M. (2002), “White blood cell count is positively correlated with albumin excretion rate in subjects with type 2 diabetes,” *Diabetes Care*, 25, 2354–2355.

Chen, X.-L., Xue, L., Wang, W., Chen, H.-N., Zhang, W.-H., Liu, K., Chen, X.-Z., Yang, K., Zhang, B., Chen, Z.-X., et al. (2015), “Prognostic significance of the combination of preoperative hemoglobin, albumin, lymphocyte and platelet in patients with gastric carcinoma: a retrospective cohort study,” *Oncotarget*, 6, 41370.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018), “Double/debiased machine learning for treatment and structural parameters,” *The Econometrics Journal*, 21, C1–C68.

Dai, D., Rigollet, P., and Zhang, T. (2012), “Deviation optimal learning using greedy Q -aggregation,” *The Annals of Statistics*, 40, 1878 – 1905.

Du, H., Siah, K. T. H., Ru-Yan, V. Z., Teh, R., Tan, C. Y. E., Yeung, W., Scaduto, C., Bolongaita, S., Cruz, M. T. K., Liu, M., et al. (2021), “Prediction of in-hospital mortality of Clostridioides difficile infection using critical care database: a big data-driven, machine learning approach,” *BMJ Open Gastroenterology*, 8, e000761.

Fialho, A. S., Cismondi, F., Vieira, S. M., Reti, S. R., Sousa, J. M., and Finkelstein, S. N. (2012), “Data mining using clinical physiology at discharge to predict ICU readmissions,” *Expert Systems with Applications*, 39, 13158–13165.

Fukui, M., Tanaka, M., Hasegawa, G., Yoshikawa, T., and Nakamura, N. (2008), “Association between serum bioavailable testosterone concentration and the ratio of glycated albumin to glycated hemoglobin in men with type 2 diabetes,” *Diabetes Care*, 31, 397–401.

Guo, Y., Shi, D., Zhang, J., Mao, S., Wang, L., Zhang, W., Zhang, Z., Jin, L., et al. (2019), “The Hemoglobin, Albumin, Lymphocyte, and Platelet (HALP) Score is a Novel Significant Prognostic Factor for Patients with Metastatic Prostate Cancer Undergoing Cytoreductive Radical Prostatectomy,” *Journal of Cancer*, 10, 81.

Huang, J., Horowitz, J. L., and Wei, F. (2010), “Variable selection in nonparametric additive models,” *Ann. Statist.*, 38, 2282–2313.

Javanmard, A. and Montanari, A. (2018), “Debiasing the lasso: Optimal sample size for Gaussian designs,” *The Annals of Statistics*, 46, 2593 – 2622.

Jhou, H.-J., Chen, P.-H., Yang, L.-Y., Chang, S.-H., and Lee, C.-H. (2021), “Plasma Anion Gap and Risk of In-Hospital Mortality in Patients with Acute Ischemic Stroke: Analysis from the MIMIC-IV Database,” *Journal of Personalized Medicine*, 11, 1004.

Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. (2016), “MIMIC-III, a freely accessible critical care database,” *Scientific Data*, 3, 160035.

Lin, Y. and Zhang, H. H. (2006), “Component selection and smoothing in multivariate nonparametric regression,” *Ann. Statist.*, 34, 2272–2297.

Liu, X., Feng, Y., Zhu, X., Shi, Y., Lin, M., Song, X., Tu, J., and Yuan, E. (2020), “Serum anion gap at admission predicts all-cause mortality in critically ill patients with cerebral infarction: evidence from the MIMIC-III database,” *Biomarkers*, 25, 725–732.

Meier, L., Van de Geer, S., and Bühlmann, P. (2009), “High-dimensional additive modeling,” *The Annals of Statistics*, 37, 3779–3821.

Nates, J. L., Nunnally, M., Kleinpell, R., Blosser, S., Goldner, J., Birriel, B., Fowler, C. S., Byrum, D., Miles, W. S., Bailey, H., et al. (2016), “ICU admission, discharge, and triage guidelines: a framework to enhance clinical operations, development of institutional policies, and further research,” *Critical Care Medicine*, 44, 1553–1602.

Nemirovski, A. (2000), “Topics in Non-Parametric Statistics,” .

Peng, D., Zhang, C.-j., Gong, Y.-q., Hao, H., Guan, B., Li, X.-s., and Zhou, L.-q. (2018), “Prognostic significance of HALP (hemoglobin, albumin, lymphocyte and platelet) in patients with bladder cancer after radical cystectomy,” *Scientific reports*, 8, 1–9.

Raskutti, G., Wainwright, M. J., and Yu, B. (2012), “Minimax-Optimal Rates For Sparse Additive Models Over Kernel Classes Via Convex Programming,” *Journal of Machine Learning Research*, 13, 389–427.

Rigollet, P. and Tsybakov, A. (2011), “Exponential Screening and optimal rates of sparse estimation,” *Ann. Statist.*, 39, 731–771.

Rigollet, P., Tsybakov, A. B., et al. (2012), “Sparse estimation by exponential weighting,” *Statistical Science*, 27, 558–575.

Rudelson, M. and Zhou, S. (2012), “Reconstruction from Anisotropic Random Measurements,” in *Proceedings of the 25th Annual Conference on Learning Theory*, eds. Mannor, S., Srebro, N., and Williamson, R. C., Edinburgh, Scotland: PMLR, vol. 23 of *Proceedings of Machine Learning Research*, pp. 10.1–10.24.

Shen, X.-B., Zhang, Y.-X., Wang, W., and Pan, Y.-Y. (2019), “The Hemoglobin, Albumin, Lymphocyte, and Platelet (HALP) Score in Patients with Small Cell Lung Cancer Before First-Line Treatment with Etoposide and Progression-Free Survival,” *Medical science monitor: international medical journal of experimental and clinical research*, 25, 5630.

Tabak, Y. P., Sun, X., Nunez, C. M., Gupta, V., and Johannes, R. S. (2017), “Predicting readmission at early hospitalization using electronic clinical data: an early readmission risk score,” *Medical Care*, 55, 267.

Tang, Y., Zeng, X., Feng, Y., Chen, Q., Liu, Z., Luo, H., Zha, L., and Yu, Z. (2021), “Association of systemic immune-inflammation index with short-term mortality of congestive heart failure: a retrospective cohort study,” *Frontiers in Cardiovascular Medicine*, 8.

Tojek, K., Kowalczyk, G., Czerniak, B., Banaś, W., Szukay, B., Korzycka-Wilińska, W., Banaszkiewicz, Z., and Budzyński, J. (2019), “Blood albumin as a prognostic factor among unselected medically treated inpatients,” *Biomarkers in medicine*, 13, 1059–1069.

Tsybakov, A. B. (2008), *Introduction to Nonparametric Estimation*, Springer Publishing Company, Incorporated, 1st ed.

Verzelen, N. (2012), “Minimax risks for sparse regressions: Ultra-high dimensional phenomena,” *Electron. J. Statist.*, 6, 38–90.

Xue, Y., Klabjan, D., and Luo, Y. (2019), “Predicting ICU readmission using grouped physiological and medication trends,” *Artificial Intelligence in Medicine*, 95, 27–37.

Zhang, Y. and Bradic, J. (2019), “High-dimensional semi-supervised learning: in search for optimal inference of the mean,” *arXiv e-prints*, arXiv:1902.00772.

Zhao, J., Leng, C., et al. (2016), “An analysis of penalized interaction models,” *Bernoulli*, 22, 1937–1961.