

HIGH-DIMENSIONAL SEMI-SUPERVISED LEARNING: IN SEARCH FOR OPTIMAL INFERENCE OF THE MEAN

YUQIAN ZHANG AND JELENA BRADIC

ABSTRACT. We provide a high-dimensional semi-supervised inference framework focused on the mean and variance of the response. Our data are comprised of an extensive set of observations regarding the covariate vectors and a much smaller set of labeled observations where we observe both the response as well as the covariates. We allow the size of the covariates to be much larger than the sample size and impose weak conditions on a statistical form of the data. We provide new estimators of the mean and variance of the response that extend some of the recent results presented in low-dimensional models. In particular, at times we will not necessitate consistent estimation of the functional form of the data. Together with estimation of the population mean and variance, we provide their asymptotic distribution and confidence intervals where we showcase gains in efficiency compared to the sample mean and variance. Our procedure, with minor modifications, is then presented to make important contributions regarding inference about average treatment effects. We also investigate the robustness of estimation and coverage and showcase widespread applicability and generality of the proposed method.

Keywords: Ridge, Lasso, Model-lean, Signal-to-noise ratio.

1. INTRODUCTION

In many applications, we want to use as much relevant data as possible to draw inferences about a population parameter of interest. For example, when estimating a causal parameter such as average treatment effect or the effect of its size, the practitioner is tempted to try to collect as much relevant data as possible. Nowadays most of the datasets are equipped with a large number of features that represent relevant genetic, medical, consumer, economic or social behavior. However, most of them have a relatively small number of responses of interest; in particular, the number of responses tends to be of one or multiple order of magnitude smaller preventing meaningfully statistical inferences about say causal parameter of interest. At the same time, there has been an explosion of many open datasets containing additional information about the features of interest.

The practitioners are therefore tempted to try to incorporate this additional information that is available to them. How to best do that, to provide inferential statistical guarantees on the parameter of interest? Statistical inference is particularly challenging in the presence of a large number of features. An impediment to providing optimal inferences is the fear that the researchers will try to “compress” the high dimensional feature space usually by employing some generic unsupervised method thus, highlighting the heterogeneity that is potentially purely spurious. The dimensionality reduction is usually made in an ad-hoc manner while keeping only sufficient statistics of the auxiliary information for example. Evidently, this may not be an optimal way of utilizing the information at

hand. In this paper, we seek to address this challenge by developing a method that will use all of the auxiliary features and therefore potentially borrowing strength across all of the available data in order to provide statistical guarantees regarding inferences about the mean and variance of the response vector. Among other applications, this work has strong implications for inference regarding the average treatment effect and the size of the treatment effect.

Approaches to semi-supervised estimation and inference regarding the response population distribution includes the work of [Zhang et al. \(2016\)](#). Although their method will perform well when the number of features is small it will quickly break down in the presence of a large number of features; after all, it is based on a least squares estimates that cannot be computed in the high-dimensional settings. In this paper, we explore the ideas of high-dimensional statistical methods to improve the performance of semi-supervised approaches in the presence of many covariates. We focus on the family of regularized algorithms but do not specify necessarily which particular regularization is used. This allows for flexible modeling in high dimensions by selecting regularization most suitable for the application in mind. Regularized methods are powerful in explaining sparse interactions among covariates, something that is especially important in environments with many covariates.

Regardless of their widespread success at prediction and model selection, there are essential difficulties that need to be cleared before regularized methods can directly be useful for semi-supervised inference. Preferably, an estimator should have a well-understood asymptotic distribution. Asymptotic normality results are especially important in inferential tasks because regularized estimates can be difficult to evaluate using, e.g., cross-validation directly. Moreover, in the setting of semi-supervised inferential tasks, it is widely recognized that additional information arising from the augmented dataset can provide better estimates when the data generating model is possibly misspecified. The asymptotics of regularized methods for misspecified models has mainly been left open, even in the case of moderate covariate dimensions.

This paper addresses these limitations, by developing a regularized method for estimation of the mean and variance of the response that allows for a tractable asymptotic distribution and therefore justifiable inference. Our proposed method is composed of simple, oracle-like bias correction step that allows for successful inference with weak assumptions on the model specifications, such as sparsity for example. We utilize sample-splitting to guarantee optimality.

We begin by developing the consistency and asymptotic normality regarding the mean of the response in the context of high-dimensional linear models. We prove these results for many variants of regularized methods each of which can be arbitrarily slow at achieving consistency in the estimation of a population l_2 projection. We then discuss the case of the variance of the response. Under the same conditions, we also develop asymptotic normality results regarding estimates of the asymptotic variance. We are then able to use these general results and show that our asymptotic normality extends to the setting of

estimation of heterogeneous treatment effects in the potential outcomes framework with unconfoundedness (Rubin, 1974).

1.1. Related Work. There has been an implicit understanding in the statistical literature that additional knowledge or data accumulated ought to lead to an improved estimation and inference: if the goal is inferential, then additional data should lead to smaller asymptotic variance at estimation. However, it is unclear if and to what extent this understanding would hold in the semi-supervised setting where we can collect additional information regarding the design of the model but not the response. In the linear regression framework, estimation of the mean and variance of the population may bring in information about the structure of the population of interest. In the presence of additional covariates, one would like to showcase an improved understanding of the population distribution by utilizing the information of the covariates. When evaluating the quality of the estimation, asymptotic distribution plays a much more critical role than in the standard prediction tasks.

The main contribution of our work is both construction of new estimates as well as an asymptotic normality results that accommodate semi-supervised as well as a high-dimensionality phenomenon. Recent results of Wasserman and Lafferty (2008); El Alaoui et al. (2016); Mai and Couillet (2017) consider the class of graph-oriented semi-supervised learning algorithms and proceed to establish estimation or prediction properties of semi-supervised estimators when the number of observed responses was much higher than the number of features in the data. New work of Chakraborty and Cai (2017); Cheng et al. (2018); Gronsbell and Cai (2018) develops semi-supervised approaches for estimation of the conditional mean of the responses. They utilize the additional information to obtain information about the marginal distribution of the covariates in order to then reduce the bias of the local linear estimator. Meanwhile, Cai and Guo (2018) proposed a semi-supervised estimator of the explained variance; apart from discussing a different question, their task requires exact model specification. Semi-supervised inference in the context of classification has had a long tradition; see Chapelle et al. (2009); Grandvalet and Bengio (2005); Zhu and Goldberg (2009); Moscovich et al. (2016). A major limitation of this line of work is that it has lacked formal statistical inference results.

A small but growing literature, including Zhang et al. (2016); Azriel et al. (2016) have considered the development of semi-supervised estimators. Much of their work is inspired by model-free regression framework; see for example Kuchibhotla et al. (2018). These papers discuss estimation of the mean and variance and use least-squares methods for which they report confidence intervals as well as asymptotic distribution. To our knowledge, however, we provide a set of conditions under which efficient estimation of the response parameters is established without firm guarantees on the model specification as well as without strong restrictions on the dimensionality of the feature space.

We, therefore, view our contribution as complementary to this growing literature, by showing that regularized methods can be effectively utilized for efficient semi-supervised inferences. The theoretical results developed here are general in that many regularized

methods are considered simultaneously. Besides, we only require mild conditions on dimensionality as well as sparsity. We believe that the new estimating tools developed here will be useful beyond the specific class of environments studied in our paper. We illustrate this point, by applying our findings to the settings of estimation and inference of heterogeneous treatment effects in the framework of potential outcomes.

Several papers use generalized methods for estimating heterogeneous treatment effects as well as treatment effect size. Random forests approaches appear in [Wager and Athey \(2017\)](#). Other related approaches include those of [Künzel et al. \(2017\)](#) and [Chernozhukov et al. \(2017\)](#), which build learners of sorts that can conform to many types of nonparametric or machine learning methods; however, these papers do not analyze semi-supervised settings and possible improvements therein over using labeled observations only. In improving treatment effects estimates over the labeled observations, we follow our work on semi-supervised estimation of the mean closely whereas for the case of treatment effect size we build further on our work on semi-supervised estimation of the variance of the responses. As we will show we can achieve asymptotic normality under relaxed assumptions, in particular, a consistent estimation is no longer needed for both of the unknowns; one consistency suffices.

1.2. Notation. With a slight abuse of notation, throughout the article a constant $c > 0$, independent of n, p, m may change value from a line to a line. A vector e_1 denotes a p -dimensional vector $(1, 0, 0, \dots, 0)^\top$. For any vector $a \in \mathbb{R}^p$, $\|a\|_2 = \sqrt{\sum_{j=1}^p a_j^2}$, $\|a\|_\infty = \max_j |a_j|$. For any matrix $A \in \mathbb{R}^{p \times p}$, we denote $\|A\|_2 = \sup_{z \neq 0} \|Az\|_2 / \|z\|_2$. Moreover, we denote $\|a\|_A^2 = a^\top A a$. For all a , we denote with a_{-j} a vector a whose j -coordinate has been deleted. For a covariate X , with the mean μ , we denote $Z = C^{-\frac{1}{2}}(X - \mu)$. Moreover, a matrix C , denotes $C = \text{Cov}X$.

1.3. Organization. In [Section 1](#), we introduce the high-dimensional semi-supervised setting and related work. In [Section 2](#), we propose the new estimators and asymptotic confidence intervals for the mean and variance of the response distribution. In [Section 3](#), we provide theoretical results pertaining to consistency as well as asymptotic normality of the newly introduced semi-supervised estimators. We discuss their benefits in comparison to sample mean and variance as well as sufficient conditions for attaining them. In [Section 4](#), we apply our semi-supervised approach to heterogeneous average treatment effects estimation as well as inference. In [Section 5](#), we illustrate the finite-sample behavior of the introduced estimators in several simulations settings and a real dataset regarding HIV drug resistance.

2. EFFICIENT ESTIMATION OF THE MEAN

Suppose that we have access to n independent and identically distributed training examples labeled $i = 1, \dots, n$, each of which consists of a covariate vector $X_i \in \mathbb{R}^{p-1}$ and a response $Y_i \in \mathbb{R}$. Semi-supervised structure of the data allows for an additional set of unlabeled observations. Namely, we assume that an additional set of m independent and

identically distributed training examples labeled $i = n + 1, \dots, n + m$ is available and consists of only covariate vectors, X_i ; specifically, we do not observe a set of responses Y_i for $i = n + 1, \dots, n + m$. We will consider the setting of $\tau := \lim_{m,n \rightarrow \infty} n/m < \infty$ where the case $\tau < 1$ is particularly useful. Although our methods apply more broadly, we will be particularly interested in a case when $p \gg n$; to the best of our knowledge even low-dimensional settings provide new results; we provide asymptotic normal result regarding the mean of the response even when conditional mean function estimates may not be asymptotically normal. Lastly, we will assume that our covariates are random. It is worth pointing that in a high-dimensional setting, a random design is a necessary step for providing inferential results.

Let $(Y, X) \sim \mathbb{P}$ denote the population response and covariate vector. We posit the existence of a vector β^* as a l_2 projection in the population space, defined as

$$(1) \quad \beta^* := \arg \min_{\beta \in \mathbb{R}^p} \mathbb{E} (Y - \beta_1 - X^\top \beta_{-1})^2,$$

where $\beta^* \in \mathbb{R}^p$ could be called the population slope. Let ε_i be the noise not explained by the linear model, i.e., $\varepsilon_i := Y_i - \beta_1^* - X_i^\top \beta_{-1}^*$. We denote with $\sigma_\varepsilon^2 := \text{Var} \varepsilon$ the unexplained variance. Under this setting, $E[\varepsilon|X] = 0$ is not required and the design and the error are possibly dependent. Consequently, this setup does not assume necessarily that the regression model is indeed a linear model.

We observe that the mean of the response, i.e., the parameter of interest that we wish to perform inferences on can be expressed now as a linear combination

$$\theta = \beta^{*\top} \tilde{\mu}.$$

The main objective is to perform efficient inferences without assuming correct model specification and also identifying what kind of specifications allow for efficient procedures.

In general, it may seem plausible to estimate θ efficiently we only need to consider a \sqrt{n} consistent estimate of $\hat{\beta}$. Although this observation holds in low-dimensional setting (Zhang et al., 2016), high-dimensional setting presents with unique challenges. Obtaining such estimator comes with a set of restricted assumptions; even if \mathbb{P}_X is known and a de-biased lasso $\hat{\beta}_{\text{db}}$ is performed, \sqrt{n} consistency of $\hat{\beta}_{\text{db}}^\top \tilde{\mu}$ is guaranteed only when $s = O(\sqrt{n}/\log p)$. Here, we propose to utilize the fact that $\tilde{\mu}$ is the mean of the covariate vectors and then propose a direct estimation of θ that allows both efficient estimation as well as $n^{-1/2}$ inference about θ under weaker set of assumptions.

2.1. Estimation of the mean in high-dimensional settings. At a high level, the parameter of interest, θ , can be seen as the linear contrast of the parameter β^* . Recent literature on inference in high-dimensional setting has illuminated the significance of knowledge of the distribution of the covariate X ; see for example Cai et al. (2017); Dicker (2014); Janson et al. (2017); Barber et al. (2015). The overwhelming conclusion is that when inference is a goal, the knowledge of the distribution of X allows for significant relaxation of the distributional assumption on $Y|X$, i.e., moving from model $Y|X$ to model X alone.

We begin by considering the simplest case of known $\tilde{\Sigma} = \mathbb{E}\tilde{X}\tilde{X}^\top$. For that setting, we know that a bias corrected estimator $\hat{\beta} + \tilde{\Sigma}^{-1}n^{-1}\sum_{i=1}^n\tilde{X}_i(Y_i - \tilde{X}_i^\top\hat{\beta})$ should be asymptotically normal for some suitable choice of $\hat{\beta}$. In this case an oracle estimator would then take the form of

$$\begin{aligned}\tilde{\mu}^\top\left(\hat{\beta} + \tilde{\Sigma}^{-1}n^{-1}\sum_{i=1}^n\tilde{X}_i(Y_i - \tilde{X}_i^\top\hat{\beta})\right) &= \tilde{\mu}^\top\hat{\beta} + n^{-1}e_1^\top\sum_{i=1}^n\tilde{X}_i(Y_i - \tilde{X}_i^\top\hat{\beta}) \\ &= \tilde{\mu}^\top\hat{\beta} + n^{-1}\sum_{i=1}^n(Y_i - \tilde{X}_i^\top\hat{\beta}),\end{aligned}$$

where due to construction, $\tilde{\Sigma}^{-1}\tilde{\mu} = e_1$ and $e_1^\top\tilde{X}_i = 1$.

In this section, we seek to build an estimate that would resemble the right hand side above. We proceed to formulate a cross-fitted formulation that requires separate sets of observations for computing the estimate of β^* and that of evaluating the size of the bias correction factor; this enables us to disentangle their dependencies better. What is remarkable is that we would show that such a construction can be stable even for leave-one-out estimates of β^* ; see Figures 5, 6.

We start by splitting the labeled observations into sets I_k each of size N (for the sake of simplicity) for $k = 1, \dots, K$. Let $\hat{\beta}^{(-k)}$ denote an estimate of β^* computed on the all but the I_k labeled observations,

$$\hat{\beta}^{(-k)} = \hat{\beta}\left(\{(\tilde{X}_i, Y_i) : i \in \{1, 2, \dots, n\} \setminus I_k\}\right) \in \mathbb{R}^p.$$

Then, it is natural to estimate θ by the cross-fitted estimate

$$(2) \quad \hat{\theta}^{(k)} = \hat{\mu}^\top\hat{\beta}^{(-k)} + N^{-1}\sum_{i \in I_k}\left(Y_i - \tilde{X}_i^\top\hat{\beta}^{(-k)}\right)$$

where an estimate of the location of X , μ , is defined on an un-labeled observations

$$(3) \quad \hat{\mu} = \frac{1}{m}\sum_{i=n+1}^{m+n}\tilde{X}_i.$$

Finally, given a procedure for generating a single estimate of θ , our method generates a family of K such estimates, each of which outputs $\hat{\theta}^{(k)}$ estimate. Our procedure then aggregates the estimates by averaging them:

$$(4) \quad \hat{\theta} = \frac{1}{K}\sum_{k=1}^K\hat{\theta}^{(k)}.$$

The advantage of the averaged estimate is that it is not clear how to best split the data to perform cross-fitting. In this case, it is better to average over potentially many estimates than to seek the best highly optimized case. One can also randomize the partition of the index set and then further average the estimate over the different generated partitions.

2.2. From the mean to the variance. Given an estimate of the first moment, classical methods would then be focused on understanding the behavior of the second moment. Often times this provides further information on the size of the effect of the first moment, for example.

A classical estimate of the variance of the response, Y , is the simple sample variance estimate

$$S_Y^2 = n^{-1} \sum_{i=1}^n (Y_i - \bar{Y})^2.$$

However, the estimate above is not utilizing any knowledge of the covariate vectors. A naive estimate that utilizes X is

$$n^{-1} \sum_{i=1}^n (Y_i - \hat{\theta})^2,$$

where $\hat{\theta}$ utilizes both the labeled as well as un-labeled observations. However, we will show that both of these estimates can be improved upon if we utilize the availability of the additional data of X as well as presence of the projection β^* . Our theoretical results regarding $\hat{\theta}$ illustrate

$$n^{-1} \sum_{i=1}^n (Y_i - \hat{\theta})^2 = \sigma_Y^2 + n^{-1} \sum_{i=1}^n \left(\beta^{*\top} \left[\tilde{V} \tilde{V}^\top - \tilde{C} \right] \beta^* \right) + T + \mathcal{O}_P(n^{-1}),$$

where T is a mean-zero component, $T = n^{-1} \sum_{i=1}^n \{ 2\beta^{*\top} \tilde{V}_i \varepsilon_i + \varepsilon_i^2 - \sigma_\varepsilon^2 \}$. Hence, to reduce the variance of $n^{-1} \sum_{i=1}^n (Y_i - \hat{\theta})^2$ we can add a correction term that utilizes the knowledge of X and β^* effectively so that the term on the right hand side above disappears. Therefore, we propose a new estimator

$$(5) \quad \hat{\sigma}_Y^{2^{(k)}} = N^{-1} \sum_{i \in I_k} (Y_i - \hat{\theta})^2 + N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \left(\hat{C} - \hat{V}_i \hat{V}_i^\top \right) \hat{\beta}^{(-k)}.$$

where the estimates of the de-meaned covariates, \tilde{V}_i , are defined as $\hat{V}_i := \tilde{X}_i - \hat{\mu}$ whereas their covariance matrix is estimated on a set of unlabeled observations

$$(6) \quad \hat{C} = \frac{1}{m} \sum_{i=n+1}^{m+n} \hat{V}_i \hat{V}_i^\top.$$

Then, a simple averaging leads to the final estimate of the variance

$$(7) \quad \hat{\sigma}_Y^2 = \frac{1}{K} \sum_{k=1}^K \hat{\sigma}_Y^{2^{(k)}}.$$

2.3. Asymptotic inference for the mean. Our results require some conditions on the estimating scheme used to obtain estimates $\hat{\beta}$. Moreover, the conditions will change depending on the goal of interest: estimation or inference. In general, the plug-in estimates need to be evaluated on a data that was not used to create terms of the bias-correction.

For estimation efficiency, the linear estimator of β^* does not have to be a consistent estimator, but the estimation efficiency will reach the parametric rate nonetheless. For efficient confidence interval constructions, the linear estimator of β^* needs to only be consistent, with the rate of consistency being significantly slower than the \sqrt{n} and yet the resulting estimator would be asymptotically normal with a parametric rate. Given these high-level conditions, we obtain a widely applicable consistency and asymptotic normality results that apply to several different algorithms used to estimate $\hat{\beta}$.

Our first result is that the proposed estimator, $\hat{\theta}$ achieves smaller mean squared error than that of \bar{Y} . Such result only requires that $\|C\|_2 > 0$ and $\hat{\beta}$ is such that $K\mathbb{E}\|\hat{\beta}^{(-k)} - \beta^*\|_2^2 \leq b^2$, for $b^2 = \lim_{n,p} \beta_{-1}^{*\top} C \beta_{-1}^*$ i.e., the explained variance. One can see that for models that are not linear but for which a linear approximation holds in such a way that the explained variance, b^2 is large, the above condition will be satisfied. Our second result is that $\hat{\theta}$ is \sqrt{n} consistent for the true mean of Y . To achieve square root consistency, we need to assume that $\|\hat{\beta} - \beta^*\|_2 = \mathcal{O}_P(1)$, as $n, p \rightarrow \infty$. In particular, $\hat{\beta}$ can be an asymptotically biased estimator; as long as that bias is not diverging, $\hat{\theta}$ will still be an asymptotically consistent estimate.

Beyond consistency, in order to do statistical inference, we need to understand the asymptotic distribution of $\hat{\theta}$. We showcase that as long as $\hat{\beta}$ is asymptotically unbiased estimate, regardless of the rate at which it achieves, consistency,

$$\sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \sigma_\varepsilon^2 + \tau b^2).$$

In particular, whenever, $\tau = 0$ as $n \rightarrow \infty$, and at least one of the covariates has positive absolute correlation with the response Y , then $\hat{\theta}$ is a more efficient estimator than \bar{Y} as it's asymptotic variance is always smaller than that of \bar{Y} ; note that a simpler requirement of $\tau < 1$ is allowed. Moreover, this result is blind to the distribution of X and ε as long as the second $2 + c$, $c > 0$ moments are bounded.

Lastly, we showcase that the asymptotic variance above can be consistently estimated and therefore confidence intervals about the true mean of Y be constructed as

$$(8) \quad \text{CI}(\theta) := \left[\hat{\theta} - z_{1-\frac{\alpha}{2}} \sqrt{\hat{\sigma}_\varepsilon^2 + \tau \hat{b}^2} / \sqrt{n}, \hat{\theta} + z_{1-\frac{\alpha}{2}} \sqrt{\hat{\sigma}_\varepsilon^2 + \tau \hat{b}^2} / \sqrt{n} \right].$$

Our results for the theoretical guarantees of $\text{CI}(\theta)$ only require consistency of $\hat{\beta}$ at an arbitrarily slow rate. To consistently estimate the variance, we estimate separately σ_ε^2 as well as b^2 . To do so, we define the cross-fitted residuals in the following form

$$\hat{\varepsilon}_i = Y_i - \hat{\theta} - \hat{\beta}^{(-k)\top} \hat{V}_i.$$

The motivation behind this careful construction, is that, given high-dimensional observations, the residual as defined above guarantees that the estimated residuals are \sqrt{n} -consistent. This, in turn, provides a more stable estimates and enables theoretically weak conditions. To see that the naive estimate $Y_i - \hat{\beta}^\top \tilde{X}_i$ may not guarantee \sqrt{n} consistency, we just need to observe that in such case $Y_i - \hat{\beta}^\top \tilde{X}_i = \varepsilon_i + (\theta - \hat{\beta}^\top \tilde{\mu}) - (\hat{\beta} - \beta^*)^\top (\tilde{X}_i - \tilde{\mu})$

and the term $(\theta - \widehat{\beta}^\top \widetilde{\mu})$ is not necessarily \sqrt{n} consistent when $p \gg n$. Our cross-fitted construction can be seen as a bias corrected estimate of the residuals.

Now we define the estimate of the variance of the noise as a simple average of the squared cross-fitted residuals

$$(9) \quad \widehat{\sigma}_\varepsilon^{2(k)} := N^{-1} \sum_{i \in I_k} \widehat{\varepsilon}_i^2.$$

To consistently estimate parameter b^2 we utilize the framework of bias-corrected estimate as developed in [Cai and Guo \(2018\)](#). We enrich the original proposal of the authors with the corrected, cross-validated residuals

$$(10) \quad \widehat{b}^{2(k)} := \widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)} + \frac{2}{N} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widehat{V}_i \widehat{\varepsilon}_i.$$

Then, the final estimate of the variance of $\widehat{\theta}$ is a simple linear combination

$$(11) \quad \widehat{\sigma}_\theta^2 = \widehat{\sigma}_\varepsilon^2 + n\widehat{b}^2/m$$

where

$$(12) \quad \widehat{\sigma}_\varepsilon^2 := K^{-1} \sum_{k=1}^K \widehat{\sigma}_\varepsilon^{2(k)}, \quad \widehat{b}^2 := K^{-1} \sum_{k=1}^K \widehat{b}^{2(k)}.$$

The choice of the partition $\mathcal{P} = \{I_k\}_{k=1}^K$ is not vital when n is large. However, we can utilize the randomness induced by sample splitting to improve the estimated variance in finite samples. Suppose we have randomly generated T partitions $\{\mathcal{P}_t\}_{t=1}^T$. Based on each partition \mathcal{P}_t , we have a semi-supervised estimator $\widehat{\theta}_t$, (4), and its estimated variance $\widehat{V}(\widehat{\theta}_t)$, (11). Then, we take the average of $\widehat{\theta}_t$'s as a new averaged estimate of θ

$$\bar{\theta} = T^{-1} \sum_{t=1}^T \widehat{\theta}_t.$$

With it we propose an estimate of its variance that includes a finite sample correction

$$\widehat{\text{Var}}(\bar{\theta}) = T^{-1} \sum_{t=1}^T \widehat{V}(\widehat{\theta}_t) + T^{-1} \sum_{t=1}^T (\widehat{\theta}_t - \bar{\theta})^2.$$

With this the new confidence intervals can be defined as

$$\text{CI}(\theta) := \left[\bar{\theta} - z_{1-\frac{\alpha}{2}} \sqrt{\widehat{\text{Var}}(\bar{\theta})/\sqrt{n}}, \bar{\theta} + z_{1-\frac{\alpha}{2}} \sqrt{\widehat{\text{Var}}(\bar{\theta})/\sqrt{n}} \right].$$

2.4. Asymptotic inference for the variance. When we are interested in estimating and perhaps constructing confidence intervals regarding the variance of Y , our procedure requires the same set of simple assumptions used in obtaining inferential statements regarding the mean of Y . Even when $\widehat{\beta}$ is a biased estimate whose bias is bounded asymptotically (but is not diminishing) we are able to guarantee \sqrt{n} consistency of the estimate,

(7). For consistent $\widehat{\beta}$ (without specified rate assumptions) we can guarantee more,

$$(13) \quad \sqrt{n}(\widehat{\sigma}_Y^2 - \sigma_Y^2) \xrightarrow{d} N\left(0, \text{Var}\left(\varepsilon^2 + 2\beta^{*\top} \widetilde{V}\varepsilon\right) + \tau \text{Var}\left(\beta^{*\top} \widetilde{V}\right)^2\right).$$

We note that the result above remains correct even when there is large dependence of ε_i on \widetilde{V}_i . The result simplifies a lot if both the covariates X_i and the errors ε_i have Gaussian distribution; in that case, the asymptotic variance becomes $2\sigma_\varepsilon^4 + 4\sigma_\varepsilon^2 b^2 + 2\tau b^4$. Moreover, under the same set of assumptions we can consistently estimate the asymptotic variance of $\widehat{\sigma}_Y^2$. To do so, we estimate the two components of the asymptotic variance separately. Let's focus on estimating $\text{Var}(\beta^{*\top} \widetilde{V})^2$ first. For that end, we construct a consistent estimates of $(\beta^{*\top} \widetilde{V})^2 - \mathbb{E}[(\beta^{*\top} \widetilde{V})^2]$, $\xi_i^{(k)}$, as follows

$$(14) \quad \xi_i^{(k)} = \widehat{\beta}^{(-k)\top} \left(\widehat{V}_i \widehat{V}_i^\top - \widehat{C} \right) \widehat{\beta}^{(-k)}.$$

Then we set

$$(15) \quad \widehat{\sigma}_\xi^2 := \frac{1}{n} \sum_{k=1}^K \sum_{i \in I_k} \xi_i^{(k)2}$$

and show that it estimates well $\text{Var}(\beta^{*\top} \widetilde{V})^2$. Next, we wish to estimate $\text{Var}\left(\varepsilon^2 + 2\beta^{*\top} \widetilde{V}\varepsilon\right)$. For that end we define

$$(16) \quad \eta_i^{(k)} = \widehat{\varepsilon}_i^2 + 2\widehat{\beta}^{(-k)\top} \widehat{V}_i \widehat{\varepsilon}_i + \widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)},$$

and observe that $\widehat{\sigma}_Y$ is an average of η_i^2 . Then, we create a cross-fitted residuals of the following form

$$(17) \quad \nu_i^{(k)} = \eta_i^{(k)} - \widehat{\sigma}_Y^2$$

and estimate $\text{Var}\left(\varepsilon^2 + 2\beta^{*\top} \widetilde{V}\varepsilon\right)$ with

$$(18) \quad \widehat{\sigma}_\nu^2 = \frac{1}{n} \sum_{k=1}^K \sum_{i \in I_k} \nu_i^{(k)2}$$

We showcase that as long as any consistent estimate of $\widehat{\beta}$ is used, the confidence intervals

$$(19) \quad \text{CI}(\sigma_Y^2) := \left[\widehat{\sigma}_Y^2 - z_{1-\frac{\alpha}{2}} \sqrt{\widehat{\sigma}_\nu^2 + \tau \widehat{\sigma}_\xi^2} / \sqrt{n}, \widehat{\sigma}_Y^2 + z_{1-\frac{\alpha}{2}} \sqrt{\widehat{\sigma}_\nu^2 + \tau \widehat{\sigma}_\xi^2} / \sqrt{n} \right]$$

will be asymptotically correct.

3. ASYMPTOTIC THEORY

Recall that β^* is defined as the l_2 projection, (1). In the interest of generality, we develop asymptotic theory of the proposed semi-supervised estimates in the context where the model relating Y to X is not necessarily linear. In this section we are still able to provide traditional \sqrt{n} asymptotics on inferring $\mathbb{E}[Y]$ and σ_Y^2 . The true conditional mean function $\mathbb{E}[Y|X]$ (regression surface) is assumed to be unknown and parameter

β^* describes the population best linear approximation of the unknown mean function (best linear approximation of the working model). Random design can still provide an asymptotically unbiased estimate of β^* ; such slope is an unbiased estimate of the average slope over the full range of the unknown true response surface. We showcase that even in this setting, efficient inference regarding the distribution of Y is still achievable in the presence of very many covariates.

Theorem 1. *Suppose that we have n independent and identically distributed samples $(Y_i, X_i) \sim \mathbb{P}$ whose marginal distributions are $(\mathbb{P}_Y, \mathbb{P}_X)$. In addition, suppose that we observe a supplementary set of m independent and identically distributed samples X_i that are drawn from the same distribution \mathbb{P}_X . Moreover, suppose that \mathbb{P}_X is such that $\sup_{\|a\|_2=1} \mathbb{E}|a^\top Z|^{2+c} < \infty$ for some $c > 0$ as well as $\|C\|_2 < C$ for some $C < \infty$ and that \mathbb{P}_Y is such that $\mathbb{E}|Y|^{2+c} < \infty$. Then, for any $K < \infty$ and any estimator $\hat{\beta}$ such that $\|\hat{\beta} - \beta^*\|_2 = o_P(1)$ as $n, p \rightarrow \infty$, semi-supervised estimator $\hat{\theta}$, (4), is asymptotically normally distributed*

$$\frac{\sqrt{n}(\hat{\theta} - \theta)}{\sigma_\varepsilon^2 + \tau b^2} \rightarrow \mathcal{N}(0, 1).$$

Moreover, the asymptotic variance can be consistently estimated using (12),

$$\hat{\sigma}_\varepsilon^2 + nb^2/m = \sigma_\varepsilon^2 + \tau b^2 + o_P(1).$$

Lastly, if in addition for some $c > 0$ and , $\mathbb{E}|Y|^{4+c} < \infty$, $\sup_{\|a\|_2=1} \mathbb{E}|a^\top Z|^{4+c} < \infty$, then

$$\frac{\sqrt{n}(\hat{\sigma}_Y^2 - \sigma_Y^2)}{\text{Var}(\varepsilon_1^2 + 2\beta^{*\top} \tilde{V}\varepsilon) + \tau \text{Var}(\beta^{*\top} \tilde{V})^2} \xrightarrow{d} \mathcal{N}(0, 1)$$

together with

$$\hat{\sigma}_\nu^2 + \frac{n}{m} \hat{\sigma}_\xi^2 = \text{Var}(\varepsilon^2 + 2\beta^{*\top} \tilde{V}\varepsilon) + \tau \text{Var}(\beta^{*\top} \tilde{V})^2 + o_P(1).$$

Moreover, the results above remain accurate regardless of the particular structure of the covariance matrix of the design. In particular, the case of correlated (or densely correlated) covariates is still permissible under the asymptotic regime above as the semi-supervised estimate is not requiring any sparsity structure on Σ^{-1} . For example, covariates with constant and large correlation are allowed. More broadly, any group or block correlation structure does not require a specifically tailored estimate $\hat{\beta}$ in order to achieve optimal inference; simple, naive, least-squares type estimates $\hat{\beta}$ suffice.

We note that Theorem 1 does not depend much on the choice of $\hat{\beta}$ as long as it is estimating the population slope sufficiently well. The condition, $\|\hat{\beta} - \beta^*\|_2 = o_P(1)$ is easily satisfied by many popular estimators. For example, Lasso, Dantzig selector, square-root Lasso (Belloni et al., 2011; Stucky and Van De Geer, 2017), elastic-net (Hebiri et al., 2011; De Mol et al., 2009) or Slope (Bogdan et al., 2015) estimator would all satisfy this condition provided a weak condition on the l_0 sparsity of the best linear approximation i.e., $s = o(n/\log(p))$. Note that Theorem 1 does not require homogeneity of the errors ε ; semi-supervised inference is robust to the heterogeneity.

The proof of this result is organized as follows. In Section 3.2 we provide results regarding the \sqrt{n} consistency of the semi-supervised estimators. Section 3.3 establishes asymptotic normality of the proposed estimators. We show that the bias decays slower than the variance and that high-dimensional accumulation of errors can be successfully ignored without requiring model selection consistency or correct model specification.

3.1. Theoretical background. There have been considerable efforts drawn to understand the impact and potential for understanding semi-supervised estimation and inference. In the context of binary classification, estimation properties through various forms of dimension reduction have been studied by, among others [Chapelle et al. \(2009\)](#); [Grandvalet and Bengio \(2005\)](#); [Zhu \(2005\)](#); [Zhu and Goldberg \(2009\)](#). For example, [Cannings et al. \(2017\)](#) recently proposed a local k-nearest neighbors which uses fewer neighbors by utilizing the additional data to estimate the marginal density of features. They provide asymptotic results for the excess risk over \mathbb{R}^p , $\mathcal{O}(n^{-4/(p+4)})$ when $\mathbb{E}\|X\|_2^{4+c} < \infty$ and $n^{2+p/r} = \mathcal{O}(m)$, where $c > 0$, $r \in (0, 2]$ is the smoothness of the covariates' density. Meanwhile, real-valued regression settings have been studied far less; see for example, [Zhu and Goldberg \(2006\)](#); [Kim et al. \(2009\)](#); [Zhou and Li \(2005\)](#); [Wasserman and Lafferty \(2008\)](#). However, to the best of our knowledge, our Theorem 1 is the first result that establishes conditions under which additional high-dimensional observations can be utilized to estimate the mean and variance of the response and provide estimates that are asymptotically normal.

Probably the closest result to ours is that of [Zhang et al. \(2016\)](#), where the authors tackle the same questions but consider observations with a finite or slowly growing number of features. They have shown that least squares estimates can be corrected by auxiliary dataset to provide normal distribution; they, require the sample size to grow at least faster than that of p^2 . When the number of features is much larger than the number of samples, the corrected least-squares estimate will not be asymptotically normal as its variance would grow too fast. Hence, their results cannot be used in high-dimensional settings of our interest.

In the high-dimensional setting, [Cai and Guo \(2018\)](#) discuss estimation of explained variance. Explained variance is primarily a property of the conditional distribution of the response with respect to the covariates. We are interested in the properties of the distribution of the responses themselves. They establish asymptotically normal results but only in the context of well-specified linear models; the error of the model is independent of the feature vector. However, the impact of semi-supervised settings is greatly amplified when the response model is not fully specified in that it represents only an approximation of the true, unknown model. Our work focuses specifically on such settings and establishes consistency and normality without relying on correct model specification.

Semi-supervised estimators have also been developed for estimating properties of the conditional distribution of the responses. For example, [Azriel et al. \(2016\)](#); [Chakraborty and Cai \(2017\)](#) proposed semi-supervised estimators for β^* . Their estimators are adaptive to model misspecification, with improved efficiencies over simple least-squares under model

misspecification. Cross-fitting seems to be extremely beneficial in the high-dimensional setting. [Rinaldo et al. \(2016\)](#) use it to develop confidence intervals for β^* and show that sample splitting increases the accuracy and robustness of inference. [Fan et al. \(2012\)](#) use it to develop consistent unexplained variance estimates in the high-dimensional setting. Recently, [Chernozhukov et al. \(2017\)](#) utilized cross-fitting ideas for developing double-robust machine learning estimates. In addition, [Cheng et al. \(2018\)](#) use cross-fitting to construct semi-supervised estimators for β^* . Cross-fitting can be traced back to the natural ideas of cross-validation. We observe that cross-fitting is especially powerful in avoiding bias accumulation and therefore avoiding the many pitfalls of high-dimensional regularized estimates.

An interesting question for the further theoretical study is that of the impact of a growing number of cross-fitted estimates. Our results discuss the fixed choice of K . Growing K , for example, at the same order of n , would result in studying leave-one-out estimates. In practice, we have observed that they lead to smaller variance. Hence, it seems plausible that under similar conditions, one can study a diverging number of sample fitting; however, that remains even to be studied in low-dimensional settings first.

3.2. \sqrt{n} consistency. We start by showing the \sqrt{n} consistency of the proposed semi-supervised estimator. Our approach relies on showing that the correction to the naive estimator allows consistency in estimation even when the estimator of the model parameter is not consistent. In order to state a formal result we need two simple conditions: one on the model, [Condition 1](#) or [2](#) and one on the candidate estimates of β^* , [Condition 3](#).

Condition 1. *Let the covariance matrix C be such that $\lambda_{\min}(C) > 0$ and $\lambda_{\max}(C) \leq c_1$ and $\sup_{\|a\|_2=1} \mathbb{E}|a^\top Z|^{2+c} < c_2$ as well as $\mathbb{E}|Y|^{2+c} < c_3$, for some bounded constants $c, c_1, c_2, c_3 > 0$.*

Condition 2. *The responses are such that $\mathbb{E}|Y|^{4+c} < c_1$ whereas the covariance matrix C satisfies, $\lambda_{\min}(C) > 0$ and $\lambda_{\max}(C) \leq c_2$ and $\sup_{\|a\|_2=1} \mathbb{E}|a^\top Z|^{4+c} < c_3$ for some bounded $c, c_1, c_2, c_3 > 0$.*

[Condition 1](#) or [2](#) are the only conditions imposed on the data-generating process and are used one at a time. Well-conditioned covariance matrix allows for a well-defined linear approximation model and the parameter β^* . A bounded variance of Y is needed for convenience of exposition only; all of the results remain to hold even if this condition is removed, however, the results would be less interpretable. [Condition 2](#) is only needed for the analysis of the estimate of the variance.

Condition 3. *$\hat{\beta}$ is an estimator for β^* that satisfies $\|\hat{\beta} - \beta^*\|_2 = \mathcal{O}_P(1)$, as $n, p \rightarrow \infty$*

This formalism allows us to consider a wide variety of candidate estimates of β^* . Flexibility is again as the estimators of β^* in the semi-supervised algorithm do not have to match the features of the data and the underlying model well. Semi-supervised estimator utilizes all the information from the unlabeled observations to improve over supervised or regression estimates. For example, regularized least-square estimators will satisfy the

above condition as long as $s = O(n/\log(p))$; for example if $n = 300$ and $p = 1000$, s can be as large as 100 in which case we know that Lasso estimate for example is not consistent.

Theorem 2. *Let Conditions 1 and 3 hold. Then, $m, n, p \rightarrow \infty$, $\hat{\theta} - \theta = \mathcal{O}_P(n^{-1/2})$. Moreover, if Condition 2 hold as well, we have $\hat{\sigma}_Y^2 - \sigma_Y^2 = \mathcal{O}_P(n^{-1/2})$.*

Theorem 2 includes simpler results related to the consistency of the introduced estimates. Namely, to guarantee $\hat{\theta} - \theta = o_P(1)$ we only require bounded $1 + c$ moments (weaker than Condition 1) together with Condition 3. Regarding, σ_Y^2 , Conditions 1 and 3 are sufficient to guarantee consistency of $\hat{\sigma}_Y^2$.

An important property of the semi-supervised estimator is that it does not rely on a strong signal-to-noise ratio to achieve \sqrt{n} consistency; for example, even if the correct model is not at all sparse, it can be shown that Lasso estimate will be equal to zero with high-probability in which case the proposed estimate will be the same as the naive \bar{Y} . Hence, there is no loss in efficiency, and it seems that the semi-supervised mean estimate is advantageous in almost all cases.

Finally, we believe that this property will extend when estimating many moments of the distribution of Y ; however, that task is beyond the scope of the current paper.

3.3. Asymptotic normality. In the previous section, we established consistency properties of the proposed estimators. In this section, we proceed to prove a stronger result; namely, that they are asymptotically normal and that they improve the efficiency of estimation by borrowing strength from the additional dataset.

3.3.1. The mean. In order to achieve asymptotic normality results, we need a slightly stronger condition on the quality of the estimate $\hat{\beta}$; it is now required to be consistent, albeit at an arbitrary rate that is potentially much smaller than $n^{-1/2}$.

Condition 4. $\hat{\beta}$ is an estimator of β^* that satisfies $\|\hat{\beta} - \beta^*\|_2 = o_P(1)$, as $n, p \rightarrow \infty$.

This assumption is satisfied for the case of Lasso estimator whenever $s = o(n/\log(p))$, i.e., for moderately sparse regimes; notice that the β^* is still just a linear approximation of the mean of $Y|X$. However, we do not require necessarily the restricted eigenvalue condition; see for example Raskutti et al. (2010); Van De Geer et al. (2009); Bühlmann et al. (2015). This assumption is also satisfied for the case of elastic-net estimates $\hat{\beta}$ as long as β^* is a sparse vector composed of mixed weak and strong signals. Observe that moderate sparsity it not sufficient to guarantee \sqrt{n} consistency of $\hat{\beta}$ and indeed our semi-supervised method does not need such a condition. Another example is that of a ridge estimator which allows for high multicollinearity among the features. Recall that, in Conditions 1 and 2, we only assumed bounded $2 + c$ and $4 + c$ moments, which allows heavy-tailed distributions for the features and noises. Adaptive Huber estimates $\hat{\beta}$ can be used in the case that X and ε do not have sub-gaussian tails (Sun et al., 2018). In general, we observe that $\hat{\beta}$ can be a biased estimate as long as the bias is decaying with n and p .

Theorem 3. *Let Conditions 1 and 4 hold. Then, as $m, n, p \rightarrow \infty$,*

$$(20) \quad \sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \sigma_\varepsilon^2 + \tau b^2),$$

provided that $\sigma_\varepsilon^2 + \tau b^2 > c$ for some constant $c > 0$.

A few comments are in order. Observe that

$$\text{Var}(\sqrt{n}\bar{Y}) = \sigma_Y^2 = \sigma_\varepsilon^2 + b^2 \geq \sigma_\varepsilon^2 + \tau b^2,$$

where we observe that the semi-supervised estimator $\hat{\theta}$ is asymptotically at least as accurate as \bar{Y} and is often more accurate, namely, the additional unlabeled data reduce the asymptotic variance by $(1 - \tau)b^2$. As τ decreases, i.e., as more unlabeled data we observe, the more accurate the proposed estimator $\hat{\theta}$ becomes. When $\tau = 0$, the asymptotic variance is equivalent to the case of known marginal distribution \mathbb{P}_X , where we see clear improvements in the size of the variance.

We illustrate next that the value of the explained variance b^2 depends on the correlation coefficient between the covariates and the responses. With the standardization $Z = C^{-\frac{1}{2}}V = C^{-\frac{1}{2}}(X - \mu)$, we have $\mathbb{E}Z = 0$ and $\text{Var}Z = \mathbb{I}_{p-1}$. Set $\rho_j := \text{Corr}(Z_j, Y)$ for each $j \in \{1, 2, \dots, p-1\}$, then

$$b^2 = \beta_{-1}^{*\top} C \beta_{-1}^* = (C^{-1} \mathbb{E}(VY))^\top C C^{-1} \mathbb{E}(VY) = \sigma_Y^2 \sum_{j=1}^{p-1} \rho_j^2.$$

We can see that the behavior of $\hat{\theta}$ depends on the sum of square of the correlations $\sum_{j=1}^{p-1} \rho_j^2$. If $\tau < 1$ and $\lim_{p \rightarrow \infty} b^2 = \lim_{p \rightarrow \infty} \sigma_Y^2 \sum_{j=1}^{p-1} \rho_j^2 > 0$, i.e., when at least one of the covariates has positive marginal correlation with the response, then the semi-supervised estimator $\hat{\theta}$ is asymptotically more accurate than \bar{Y} .

Theorem 4. *Let Conditions 1 and 4 hold. Then, for $\hat{\sigma}_\varepsilon^2$ and \hat{b}^2 defined in (12) we have $\hat{\sigma}_\varepsilon^2 = \sigma_\varepsilon^2 + o_P(1)$ as well as $\hat{b}^2 = b^2 + o_P(1)$, whenever $m, n, p \rightarrow \infty$.*

It is worth noting that many estimators can provide consistency of b^2 and σ_ε^2 ; simple plug-in estimates would suffice whenever estimates of β^* are consistent. However, we will show that the estimates above have smaller error rate. To illustrate that, let $\|\mathbf{x}^*\|_{\tilde{C}}^2 = \mathbf{x}^{*\top} \tilde{C} \mathbf{x}^*$. Under the Condition 2 we showcase that

$$\hat{b}^2 = T_1 + \mathcal{O}_P \left(\|\hat{\beta} - \beta^*\|_{\tilde{C}}^2 + n^{-\frac{1}{2}} \|\hat{\beta} - \beta^*\|_{\tilde{C}} + n^{-1} \right)$$

with T_1 being the term that is asymptotically normal; $T_1 = 2n^{-1} \sum_{i=1}^n \varepsilon_i \beta^{*\top} \tilde{V}_i + m^{-1} \sum_{i=n+1}^{m+n} (\beta^{*\top} \tilde{V}_i)^2$. Under the same conditions, a plug-in estimate of b^2 , $\hat{\beta}^{(-k)\top} \hat{C} \hat{\beta}^{(-k)}$, is equal to b^2 up to the order of $\mathcal{O}_P \left(\|\hat{\beta} - \beta^*\|_{\tilde{C}} \right)$.

Based on Theorems 3 and 4 it is not difficult to see that the constructed confidence interval for $\mathbb{E}[Y]$, $\text{CI}(\theta)$ as presented in (8) is asymptotically more accurate in the sense of having smaller width asymptotically, than the interval that is solely based on $\{Y_i\}_{i=1}^n$

$$[\bar{Y} - z_{1-\alpha/2} S_Y / \sqrt{n}, \bar{Y} + z_{1-\alpha/2} S_Y / \sqrt{n}].$$

This in turn, implies that the proposed semi-supervised estimator, $\hat{\theta}$ is asymptotically more accurate than the sample mean, \bar{Y} .

Moreover, the expected length of CI(θ), (8), is of the order

$$\mathcal{O}\left(\left(\frac{1}{\sqrt{n}} + \frac{1}{\sqrt{m}}\right)\left(1 + \frac{s \log p}{n}\right)\right)$$

whenever $\hat{\beta}$ is a lasso estimate and β^* is an s sparse vector. Cai et al. (2017) proved that the minimax expected length of the confidence intervals for $\theta = \xi^\top \beta^*$ over a s -sparse parameter space has the order $\mathcal{O}(1/\sqrt{n} + s \log p/n)$ when C and σ_ε are unknown. When both C and σ_ε^2 are known, they showed that the order should be $\mathcal{O}(1/\sqrt{n})$. In our setting, σ_ε^2 is unknown, but we have additional information for the covariance matrix C from the unlabeled observations. As a result, we can see that the length of our proposed confidence interval for θ has the order $\mathcal{O}_P(1/\sqrt{n})$ when $m \geq n$ and in that sense it reaches optimal length.

Throughout the paper we mainly focus on the case of the signal to noise ratio, $\text{snr} = b^2/\sigma_\varepsilon^2$, being bounded away from 0 and ∞ . However, observe that the two extremes are not particularly informative. Namely, the case of $\text{snr} = 0$ illustrates that no estimator can improve the naive \bar{Y} . Conversely, the case of $\text{snr} = \infty$ and $\tau = 0$, illustrates that the semi-supervised estimator can potentially lead to a better than \sqrt{n} convergence rate; an example includes a Lasso estimate with $\hat{\sigma}_\varepsilon = \mathcal{O}_P(1/\sqrt{n})$ that leads to the rate of $\mathcal{O}_P(\sqrt{s \log(p)}/n)$. If l_2 estimation error of $\hat{\beta}$ is of a smaller order than σ_ε then the confidence interval is still valid. However, whenever such l_2 estimation error is at least as the same order of σ_ε , one can construct new outcomes whose variance will match that of the estimation error and use those pseudo-outcomes to construct valid confidence sets.

3.3.2. *The variance.* Our analysis of the variance estimate, $\hat{\sigma}_Y$ builds upon establishing the following asymptotic representation

$$\begin{aligned} \sqrt{n}(\hat{\sigma}_Y^2 - \sigma_Y^2) &= \frac{\sqrt{n}}{m} \sum_{i=n+1}^{m+n} \beta^{*\top} \left(\tilde{V}_i \tilde{V}_i^\top - \tilde{C} \right) \beta^* + \frac{1}{\sqrt{n}} \sum_{i=1}^n (\varepsilon_i^2 - \sigma_\varepsilon^2) \\ &\quad + \frac{2}{\sqrt{n}} \sum_{i=1}^n \beta^{*\top} \tilde{V}_i \varepsilon_i + \mathcal{O}_P\left(\frac{\max_k \|\hat{\beta}^{(-k)} - \beta^*\|_2}{\sqrt{n}} + \frac{1}{n}\right). \end{aligned}$$

Observe that, asymptotic distribution of the variance, does depend on the fourth-order moments, hence the condition above does not seem unreasonable.

Theorem 5. *Let Conditions 2 and 4 hold Then, as $m, n, p \rightarrow \infty$,*

$$(21) \quad \sqrt{n}(\hat{\sigma}_Y^2 - \sigma_Y^2) \xrightarrow{d} \mathcal{N}\left(0, \text{Var}(\varepsilon^2 + 2\beta^{*\top} \tilde{V} \varepsilon) + \tau \text{Var}(\beta^{*\top} \tilde{V})\right),$$

provided that $\text{Var}(\varepsilon^2 + 2\beta^{*\top}\tilde{V}\varepsilon) + \tau\text{Var}(\beta^{*\top}\tilde{V})^2 > c$ for some constant $c > 0$. Moreover, for $\hat{\sigma}_\nu$ and $\hat{\sigma}_\xi$ defined in (15) and (18), respectively, we have

$$(22) \quad \hat{\sigma}_\nu^2 + \frac{n}{m}\hat{\sigma}_\xi^2 = \text{Var}(\varepsilon^2 + 2\beta^{*\top}\tilde{V}\varepsilon) + \tau\text{Var}(\beta^{*\top}\tilde{V})^2 + o_P(1).$$

Based on Theorem 5, when the data follows Gaussian distribution, it is not difficult to see that

$$\text{Var}(\sqrt{n}(\hat{\sigma}_Y^2 - \sigma_Y^2)) \leq \text{Var}(Y_1 - \theta)^2 = \text{Var}(\sqrt{n}S_Y^2) - o(1).$$

Namely, the constructed confidence interval for σ_Y^2 as presented in (19) is asymptotically more accurate in the sense of having smaller width asymptotically, than the interval that is solely based on $\{Y_i\}_{i=1}^n$

$$\left[S_Y^2 / \left(1 - z_{\alpha/2} \sqrt{\hat{\gamma} - 1} / \sqrt{n}\right), S_Y^2 / \left(1 + z_{\alpha/2} \sqrt{\hat{\gamma} - 1} / \sqrt{n}\right) \right],$$

or its robust alternatives (see for example Hummel et al. (2005)) where $\hat{\gamma}$ is any consistent estimator for the kurtosis. One of the choices for $\hat{\gamma}$ can be

$$\hat{\gamma} = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_i \frac{(Y_i - \bar{Y})^4}{S_Y^4} - \frac{3(n-1)^2}{(n-2)(n-3)} + 3.$$

This in turn, implies that the proposed semi-supervised estimator, $\hat{\sigma}_Y^2$ is asymptotically more accurate than the sample variance, S_Y^2 .

A general efficiency of the proposed semi-supervised estimator $\hat{\sigma}_Y^2$ would depend on the particular model of non-linearity i.e., on the particular deviations from the linear model. We illustrate the discussion with two specific examples. For that end, we introduce a proportionality coefficient r as the proportion of the decrease achieved by the semi-supervised estimator compared to S_Y^2 . We define such coefficient with

$$(23) \quad r := \frac{\text{Var}(Y - \theta)^2 - \text{Var}(\varepsilon^2 + 2\beta^{*\top}\tilde{V}\varepsilon)}{\text{Var}(Y - \theta)^2}.$$

Here, we have assumed that $m \gg n$ and the effect of τ is negligible.

The first example, discusses heteroscedastic linear model where the variance of the error depends quadratically on the covariates. The second, discusses a larger deviations of normality, where the response model is highly non-linear. In particular, we consider

$$Y_i = \sum_{j=1}^p X_{ij} + \left(a \sum_{j=1}^p X_{ij}^2 + \sum_{j=1}^p X_{ij} \right) \eta_i \quad (\text{Example 1})$$

$$Y_i = a \left| \log \left(0.8 \left| \sum_{j=1}^p X_{ij} \right| + 0.01 \right) \right| + \sum_{j=1}^p X_{ij} + \eta_i \quad (\text{Example 2})$$

where a measures the size of the deviation from the linear model. In the above $X_i, \eta_i \sim \mathcal{N}(0, 1)$. When $r > 0$ we see that the proposed estimator is more efficient than S_Y^2 .

From Example 1 we observe precisely that. The result persists over a broad range of heteroscedastic model specifications. We also observe that the larger the magnitude of




FIGURE 1. Proportion of the decrease in asymptotic variance achieved by the proposed semi-supervised estimator $\hat{\sigma}_Y^2$ as a function of the coefficient a representing the heteroscedasticity in (a) and (c) and non-linearity in (b) and (d) subfigures. Different colors correspond to different dimensionality setting of the Examples 1 and 2.

a is – the more significant the effect of heteroscedasticity is – the smaller the signal is in the linear model. This, in turns, results in smaller r values. Example 2 is showing a more complex scenario; smaller magnitudes of a indicate not too great deviations from the linear model and result in greater efficiency in $\hat{\sigma}_Y^2$. For larger a , the linear approximation is too far from the data generating process for a linear semi-supervised estimator to make contributions. Results are presented in Figure 1 where we also showcase the Signal to Noise ratio corresponding to each setting.

4. HETEROGENEOUS TREATMENT EFFECTS

Suppose we have access to n independent and identically distributed training examples, labeled $i = 1, \dots, n$, each of which consists of a feature vector $X_i \in \mathbb{R}^p$, a response $Y_i \in \mathbb{R}$ and a treatment indicator $D_i \in \{0, 1\}$. In addition to that we assume that we have access to m additional unlabeled training examples, each of which consists of a

feature vector $X_i \in \mathbb{R}^p$, $i = n + 1, \dots, n + m$ as well as treatment indicator $D_i \in \{0, 1\}$. Following the potential outcomes framework, (Neyman, 1923; Rubin, 1974; Holland, 1988) we then hypothesize the presence of potential outcomes $Y_i(0)$ and $Y_i(1)$ corresponding to respectively to the response the i -th subject would have experienced with and without the treatment; however, we only get to observe one of the them for each i . Our goal is to estimate and develop asymptotic normality for the average causal treatment effect from observational data, adjusting appropriately for confounding.

Similarly as in Section 2, we hypothesize the existence of the l_2 slopes

$$(24) \quad \beta_w^* := \min_{\beta \in \mathbb{R}^p} \mathbb{E}(Y_i(w) - \tilde{X}_i^\top \beta)^2,$$

defined at the population level for $w \in \{0, 1\}$. The response $Y_i(w)$ is the ‘‘factual outcome’’ which we observe in the data, whereas $Y_i(1-w)$ is the unrealized ‘‘counterfactual outcome.’’ We then observe that the average treatment effect takes the linear form

$$(25) \quad \delta = (\beta_1^* - \beta_0^*)^\top \tilde{\mu} := \tau_1 - \tau_0.$$

A standard way for construction of estimates related to the average treatment effects is to posit a model on the treatment assignment and then adjust for possible confounding. Treatments are assigned to subjects according to an underlying scheme that depends on the subjects’ features, i.e. $D_i \not\perp X_i$. Their dependence can be captured by the following model

$$(26) \quad D_i = e(X_i) + \zeta_i$$

where $e(X_i)$, an unknown function, is typically regarded as the propensity score (Rosenbaum and Rubin, 1983). In the following we assume two primitive conditions. The first, widely regarded as an overlap condition, i.e., that the propensity score be uniformly bounded away from 0 and 1 for all values in the support of the covariates. The second, can be thought as an identifiability condition.

Condition 5. Let $\mathbb{P}(c \leq e(X) \leq 1 - c) = 1$ with some constant $c \in (0, 1)$.

Condition 6. Denote $\varepsilon_i := Y_i(D_i) - (D_i\beta_1^* + (1-D_i)\beta_0^*)^\top \tilde{X}_i$. Let $\mathbb{E}(\varepsilon|X, D) = 0$, $\mathbb{E}(\zeta|X) = 0$, as well as $0 < \sigma_\varepsilon < \infty$.

In this section, similar to before, we would like to consider a large family of estimates of the unknowns. Therefore, $\hat{\beta}_1$, $\hat{\beta}_0$ and \hat{e} are estimators for β_1^* , β_0^* and e , respectively, which satisfy

$$(27) \quad \left(\mathbb{E}_{P_X} \left((\hat{\beta}_w - \beta_w^*)^\top \tilde{X} \right)^2 \right)^{\frac{1}{2}} = O_P(a_{n,p}), \text{ and } \left(\mathbb{E}_{P_X} \left(\hat{e}(X) - e(X) \right)^2 \right)^{\frac{1}{2}} = O_P(b_{m,p}),$$

where P_X is the marginal distribution of X , and \mathbb{E}_{P_X} denotes the conditional expectation with respect to P_X . Here, $a_{n,p}$ and $b_{m,p}$ are two non-negative sequences of numbers.

4.1. Inference regarding average treatment effects. We explore the representation (25) when constructing an efficient estimate of the ATE, δ . As our framework above

adjusts for the presence of possible confounding, the semi-supervised estimator 2 also needs to be adjusted for the confounding effects. For that end, we introduce

$$(28) \quad \widehat{\theta}^{(k)} = \widehat{\mu}^\top \widehat{\beta}^{(-k)} + N^{-1} \sum_{i \in I_k} \omega_i \left(Y_i - \widetilde{X}_i^\top \widehat{\beta}^{(-k)} \right)$$

In the above the weights, ω_i will be corresponding to the ratio of the observed treatment proportion; then, the framework from Section 2.1 will still lead to \sqrt{n} consistent estimates. We denote these weights for the treated and control, respectively, as

$$r_i^{(-k)} = D_i / \widehat{e}^{(-k)}(X_i), \quad \rho_i^{(-k)} = (1 - D_i) / (1 - \widehat{e}^{(-k)}(X_i))$$

and propose weighted semi-supervised estimates of the mean effects in the treated and control group as

$$(29) \quad \widehat{\tau}_1^{(k)} = \widehat{\beta}_1^{(-k)\top} \widehat{\mu} + N^{-1} \sum_{i \in I_k} r_i^{(-k)} (Y_i - \widehat{\beta}_1^{(-k)\top} \widetilde{X}_i), \quad \widehat{\tau}_0^{(k)} = \widehat{\beta}_0^{(-k)\top} \widehat{\mu} + N^{-1} \sum_{i \in I_k} \rho_i^{(-k)} (Y_i - \widehat{\beta}_0^{(-k)\top} \widetilde{X}_i).$$

Note that $\widehat{\mu}$ is defined much in the same way throughout the paper. Then the estimate of the ATE can be simply defined as a difference of the two estimates

$$(30) \quad \widehat{\delta}^{(k)} = \widehat{\tau}_1^{(k)} - \widehat{\tau}_0^{(k)}.$$

Of course since the choice of the sample split is arbitrary, we average the estimates

$$(31) \quad \widehat{\delta} = K^{-1} \sum_{k=1}^K \widehat{\delta}^{(k)}.$$

With the above we are ready to present the main result of this section.

Theorem 6. *Let Conditions 1, 5 and 6 hold. Then, under the setting of this section*

$$\widehat{\delta} - \delta = \mathcal{O}_P \left(n^{-\frac{1}{2}} + a_{n,p} b_{m,p} \right),$$

whenever $a_{n,p} = \mathcal{O}(1)$, $b_{m,p} \leq c_1$ for a constant $c_1 < c \in (0, 1)$, where c is the constant in Condition 5. Therefore, whenever $a_{n,p} b_{m,p} = o(1)$, $\widehat{\delta}$ is consistent. If however, $a_{n,p} b_{m,p} = \mathcal{O}(n^{-\frac{1}{2}})$, $\widehat{\delta}$ is a \sqrt{n} -consistent estimate of δ . Additionally, the asymptotic normality follows

$$\sqrt{n}(\widehat{\delta} - \delta) \xrightarrow{d} N(0, V_\delta)$$

whenever,

$$(32) \quad a_{n,p} = o(1), \quad b_{m,p} = o(1), \quad a_{n,p} b_{m,p} = o(n^{-\frac{1}{2}}),$$

with an asymptotic variance

$$(33) \quad V_\delta = \text{Var} \left(\frac{\varepsilon \zeta}{e(X)(1 - e(X))} \right) + \tau(\beta_1^* - \beta_0^*)^\top \widetilde{C}(\beta_1^* - \beta_0^*).$$

A few remarks are in order. Suppose the sparsity of the outcome model and the treatment model are s_Y and s_D , respectively. For illustration purposes suppose that both models are parametric and linear. Then, rates $a_{n,p}, b_{m,p}$ can be computed for a Lasso

regularized estimates: $a_{n,p} = \mathcal{O}_P(\sqrt{s_Y \log p/n})$, $b_{m,p} = \mathcal{O}_P(\sqrt{s_D \log p/m})$ are required to achieve asymptotic normality. Therefore, according to (32), $s_Y = o(n/\log p)$, $s_D = o(m/\log p)$ and $s_Y s_D = o(m/(\log p)^2)$. We can see that, when m is large enough that $s_D n \log p/m \rightarrow 0$, the sparsity assumption we need for the outcome model is only $s_Y = o(n/\log p)$, which is extremely mild. For large enough m this would amount to only requiring consistency, at any arbitrary rate, of the propensity model. Note that most high-dimensional inferential tools require stronger sparsity assumptions and usually the rate of $o(\sqrt{n}/\log p)$ for both the outcome as well as the propensity model. At the same time, we achieve a more efficient estimator, i.e. reduction in the size of the asymptotic variance (Chernozhukov et al., 2018).

Besides, if one is only interested in obtaining a \sqrt{n} consistent ATE estimator, as long as $m/n \rightarrow 0$, surprisingly our conditions do not require consistent estimation of the nuisance parameters – both for the outcome as well as the propensity model. This is a significant reduction of conditions compared to the existing literature; where one or the other or both of them need to be at least consistently estimated. In the context of high-dimensional or complex models, provided we have enough additional covariate information, our method allows for the outcome or propensity to be completely dense models or to be estimated using many of the machine learning methods, such as random forests, Bayesian CART or deep neural networks for example.

4.1.1. *Construction of confidence sets.* We provide a consistent estimate of the asymptotic variance V_δ with the purpose of constructing confidence sets regarding δ . For that purpose we represent V_δ as a sum of two terms; corresponding to (33), $V_\delta = V_1 + \tau V_2$.

We discuss the second component first. Observe that $V_2 = (\beta_1^* - \beta_0^*)^\top \mathbb{E}(\tilde{X}_i - \tilde{\mu})(\tilde{X}_i - \tilde{\mu})^\top (\beta_1^* - \beta_0^*) = \mathbb{E}((\beta_1^* - \beta_0^*)^\top (\tilde{X}_i - \tilde{\mu}))^2$. Hence, we propose

$$\widehat{V}_2^{(k)} := N^{-1} \sum_{i \in I_k} \xi_{\delta,i}^2$$

where

$$(34) \quad \xi_{\delta,i} := \left(\widehat{\beta}_1^{(-k)} - \widehat{\beta}_0^{(-k)} \right)^\top (\tilde{X}_i - \widehat{\mu}).$$

Here, we have evaluated the size of the vectors $\xi_{\delta,i}$ on an out-of-sample to compute $\widehat{V}_2^{(k)}$. Next, we consider the first term, where we observe

$$V_1 = \text{Var}(\varepsilon \zeta / e(X)(1 - e(X))) = \text{Var} \left(r(Y - \beta_1^* \tilde{X}) - \rho(Y - \beta_0^* \tilde{X}) \right).$$

Hence a natural estimate is then constructed by

$$\widehat{V}_1^{(k)} := N^{-1} \sum_{i \in I_k} \nu_{\delta,i}^2$$

where $e_{\delta,i} := \widehat{\delta} - (\widehat{\beta}_1^{(-k)} - \widehat{\beta}_0^{(-k)})^\top \widehat{\mu}$

$$(35) \quad \nu_{\delta,i} := r_i^{(-k)}(Y_i - \widehat{\beta}_1^{(-k)\top} \tilde{X}_i) - \rho_i^{(-k)}(Y_i - \widehat{\beta}_0^{(-k)\top} \tilde{X}_i) - e_{\delta,i}$$

Finally, our proposed estimator for V_δ is

$$(36) \quad \widehat{V}_\delta := K^{-1} \sum_{k=1}^K \left\{ \widehat{V}_1^{(k)} + \frac{n}{m} \widehat{V}_2^{(k)} \right\}.$$

Theorem 7. *Assume the same setting as in Theorem 6. Then, as $m, n, p \rightarrow \infty$,*

$$(37) \quad \widehat{V}_\delta = V_\delta + o_P(1).$$

Hence, an asymptotic $(1 - \alpha)$ -level confidence interval for the ATE δ could be defined as:

$$(38) \quad \left[\widehat{\delta} - z_{1-\frac{\alpha}{2}} \sqrt{\widehat{V}_\delta / \sqrt{n}}, \widehat{\delta} + z_{1-\frac{\alpha}{2}} \sqrt{\widehat{V}_\delta / \sqrt{n}} \right].$$

4.2. Signal-to-noise ratio. We now turn to estimating and constructing inference regarding the signal-to-noise ratio, defined as

$$d := \delta / \sigma$$

where $\sigma^2 = \text{Var}Y$. This quantity is also related to the treatment effect size and can be interpreted as what fraction of the outcome is explained by the average treatment effect. We aim to provide a consistent and asymptotically normal estimate for which we can also construct confidence intervals. Let $\widehat{\delta}$ be defined as (31). We next discuss the estimate of σ . We take on the perspective similar to Section 2.2 with minor clarifications. Notice that for each i ,

$$(39) \quad Y_i = \beta_1^{*\top} D_i \widetilde{X}_i + \beta_0^{*\top} (1 - D_i) \widetilde{X}_i + \varepsilon_i,$$

therefore for $\widetilde{W}_i = (D_i \widetilde{X}_i^\top, (1 - D_i) \widetilde{X}_i^\top)^\top \in \mathbb{R}^{2p}$ we can define the population l_2 projection

$$\beta_{\widetilde{W}}^* = \arg \min_{\beta \in \mathbb{R}^{2p}} \mathbb{E}(Y_1 - \widetilde{W}_1^\top \beta)^2$$

as the population slope. Then, from Section 2.2, we have a semi-supervised estimator $\widehat{\sigma}^2$ defined as in (7). Our proposed estimator for d is then defined as

$$(40) \quad \widehat{d} := \frac{\widehat{\delta}}{\widehat{\sigma}} \mathbb{I}\{\widehat{\sigma} \neq 0\}.$$

Here, we set \widehat{d} to be zero when $\widehat{\sigma} = 0$ in order to make \widehat{d} well defined. Notice that $\widehat{\sigma}$ is a \sqrt{n} consistent estimator under the assumptions in Theorem 8 and σ is assumed to be positive; hence, we can easily argue $\mathbb{P}(\widehat{\sigma} = 0) = o(1)$.

Theorem 8. *Assume the same setting as in Theorem 6. Moreover, assume Condition 2 holds. Then, as $m, n, p \rightarrow \infty$,*

$$(41) \quad \sqrt{n}(\widehat{d} - d) \xrightarrow{d} N(0, V_d),$$

where

$$(42) \quad V_d = \text{Var} \left(\frac{\nu_{\delta,i}^*}{\sigma} - \frac{\delta \nu_i^*}{2\sigma^3} \right) + \tau \text{Var} \left(\frac{\xi_{\delta,i}^*}{\sigma} - \frac{\delta \xi_i^*}{2\sigma^3} \right) := V_3 + \tau V_4.$$

In the above display we used the following notation

$$(43) \quad \nu_i^* = \varepsilon_i^2 - \sigma_\varepsilon^2 + 2\beta_{\widetilde{W}}^{*\top} \widetilde{V}_i \varepsilon, \quad \xi_i^* = (\beta_{\widetilde{W}}^{*\top} \widetilde{V}_i)^2 - b^2,$$

$$(44) \quad \nu_{\delta,i}^* = \frac{D_i \varepsilon_i}{e(X_i)} - \frac{(1 - D_i) \varepsilon_i}{1 - e(X_i)}, \quad \xi_{\delta,i}^* = (\beta_1^* - \beta_0^*)^\top \widetilde{V}_i.$$

To construct confidence intervals, we next provide an estimate of the asymptotic variance V_d . Let $\widehat{\sigma}^2$ and $\widehat{\delta}$ denote the semi-supervised estimators as in (7) and (30). For each $i \in I_k$, set

$$(45) \quad \nu_{d,i} := \frac{1}{\widehat{\sigma}} \nu_{\delta,i} - \frac{\widehat{\delta}}{2\widehat{\sigma}^3} \nu_i, \quad \xi_{d,i} := \frac{1}{\widehat{\sigma}} \xi_{\delta,i} - \frac{\widehat{\delta}}{2\widehat{\sigma}^3} \xi_i,$$

where ν_i and ξ_i are defined as

$$(46) \quad \nu_i = \widehat{\beta}_{\widetilde{W}}^{(-k)\top} \widehat{C}_{\widetilde{W}} \widehat{\beta}^{(-k)} + (Y_i - \widehat{\theta})^2 - (\widehat{\beta}_{\widetilde{W}}^{(-k)\top} (\widetilde{W}_i - \widehat{\mu}_{\widetilde{W}}))^2 - \widehat{\sigma}^2,$$

$$(47) \quad \xi_i = (\widehat{\beta}_{\widetilde{W}}^{(-k)\top} (\widetilde{W}_i - \widehat{\mu}_{\widetilde{W}}))^2 - \widehat{\beta}_{\widetilde{W}}^{(-k)\top} \widehat{C}_{\widetilde{W}} \widehat{\beta}^{(-k)},$$

similarly as in (17) and (14) and $\widehat{\sigma}$ is the semi-supervised estimate of $\text{Var}Y$ as mentioned before (40). Here, $\widehat{\theta}$ is the semi-supervised estimator for $\theta = \mathbb{E}Y$ proposed in (4). In addition, $\widehat{\mu}_{\widetilde{W}}$ and $\widehat{C}_{\widetilde{W}}$ are the sample mean and the sample variance of $\{\widetilde{W}_i\}_{i=n+1}^{m+n}$, respectively. Besides, $\nu_{\delta,i}$ and $\xi_{\delta,i}$ are defined as in (35) and (34), respectively. Let

$$(48) \quad \widehat{V}_3^{(k)} := N^{-1} \sum_{i \in I_k} \nu_{d,i}^2, \quad \widehat{V}_4^{(k)} := N^{-1} \sum_{i \in I_k} \xi_{d,i}^2,$$

then $\widehat{V}_3^{(k)}$ and $\widehat{V}_4^{(k)}$ are estimators for $\text{Var}(\nu_{\delta,i}^*/\sigma - \delta\nu_i^*/(2\sigma^3))$ and $\text{Var}(\xi_{\delta,i}^*/\sigma - \delta\xi_i^*/(2\sigma^3))$, respectively. Final estimates are obtained by averaging over k . Our estimator for V_d is defined as

$$(49) \quad \widehat{V}_d := \widehat{V}_3 + \frac{n}{m} \widehat{V}_4.$$

The next result showcases consistency of such estimate

Theorem 9. *Assume the same setting as in Theorem 8. Then, as $m, n, p \rightarrow \infty$,*

$$\widehat{V}_d = V_d + o_P(1).$$

Now we are able to construct an asymptotic $(1 - \alpha)$ -level confidence interval for d

$$(50) \quad \left[\widehat{d} - z_{1-\frac{\alpha}{2}} \sqrt{\widehat{V}_d}/\sqrt{n}, \widehat{d} + z_{1-\alpha/2} \sqrt{\widehat{V}_d}/\sqrt{n} \right].$$

5. FINITE-SAMPLE EXPERIMENTS

In this section, we showcase the finite-sample estimation properties of the proposed estimates. Numerical experiments consist of two different model settings: linear and non-linear. We focus on somewhat difficult cases where even if the model is linear, the inverse of the covariance or covariance matrix itself lacks strong sparsity structures. Sparsity in covariance leads to good prediction properties of Lasso estimates, whereas sparse inverse

covariance matrix leads to good inferential results. We avoid them both and yet observe the excellent performance of the proposed estimates.

5.1. Numerical experiments. Model 5.1. We first consider a setting that X_1 and ε_1 are both Gaussian and the linear model is correct. However, the model does not have sparse feature correlations or sparse precision matrix. We generate

$$X_1, X_2, \dots, X_{m+n} \sim^{\text{iid}} N_{p-1}(1_{p-1}, C_1),$$

where C_1 is a matrix that is not quite sparse. We generate the covariance matrix C_1 as follows:

$$(51) \quad D_1 := \begin{bmatrix} 0.1 & & & & \\ & 0.1 + \frac{1.9}{p-2} & & & \\ & & \ddots & & \\ & & & 0.1 + \frac{1.9(p-3)}{p-2} & \\ & & & & 2 \end{bmatrix}, \quad A := \begin{bmatrix} 1 & 0.8 & \dots & 0.8 \\ 0.8 & 1 & \dots & 0.8 \\ \vdots & \vdots & \ddots & \vdots \\ 0.8 & 0.8 & \dots & 1 \end{bmatrix}.$$

By eigen-decomposition, $A = U\Lambda U^\top$, where U is an orthogonal matrix. Define the covariance matrix as $C_1 := UD_1U^\top$. The matrix C_1 has the smallest eigenvalue $\lambda_{\min}(C_1) = 0.1$ and the largest eigenvalue $\lambda_{\max}(C_1) = 2$. We plot the heatmaps for the first 50 rows and 50 columns of the Correlation matrices converted by C_1 and C_1^{-1} in Figure 2.

We can see that neither of C_1 and C_1^{-1} are sparse. Now, we generate the n labels:

$$Y_i = X_i^\top \beta^0 + \delta_i, \quad i = 1, 2, \dots, n,$$

where $\delta_1, \delta_2, \dots, \delta_n \sim^{\text{iid}} \mathcal{N}(0, 1)$ and $\beta^0 = (\beta_0 1_{s_0}^\top, 0^\top)^\top \in \mathbb{R}^{p-1}$. As claimed before, we consider the case that the signal-to noise ratio is bounded away from 0 and ∞ . Hence, for a fixed $\sigma_\varepsilon^2 > 0$ and \mathbb{P}_X , as the sparsity $s = \|\beta^*\|_0$ grows, the individual signals $(|\beta_j^*|)_{j \in \{\beta_j^* \neq 0\}}$ should decrease so that SNR is bounded away from ∞ . β_0 is a constant such that $\text{Var}(X_1^\top \beta^0) = 1$, so the signal-to-noise ratio $b^2/\sigma_\varepsilon^2 = 1$. Fix $K = 2$, set $\widehat{\beta}^{(-k)}$ as the square-root Lasso estimator with tuning parameter $\lambda = \sqrt{\log p/(n - N)}$. Let $\widehat{\theta}$ and $\widehat{\sigma}_Y^2$ be the semi-supervised estimators for $\theta = \mathbb{E}Y_1$ and $\sigma_Y^2 = \text{Var}Y_1$ defined as in (4) and (7), respectively. Then, we propose the confidence intervals for θ and σ_Y^2 as in (8) and (19).

We repeat the simulation for 500 times. Table 1 shows the mean squared error (MSE), the average length (AL) and the average coverage (AC) of confidence intervals for θ estimated by \bar{Y} and $\widehat{\theta}$. Table 2 shows the mean squared error (MSE), the average length (AL) and the average coverage (AC) of confidence intervals for σ_Y^2 estimated by S_Y^2 and $\widehat{\sigma}_Y^2$. Lastly, Figures 3a and 3b show the MSE of \bar{Y} , $\widehat{\theta}$, S_Y^2 and $\widehat{\sigma}_Y^2$ as n varies for several choices of the sparsity s_0 .

All of the results confirm efficiency of the proposed estimates. Shorter length of confidence intervals as well as better coverage can be seen in both of Tables 1 and 2. It should be noted, that setting of $s_0 = 15$ while $n = 100$ does not satisfy outcome model sparsity conditions: however, if interested in estimation of the mean, our estimates are



(A) The heatmap of $(\text{cov2cor}(C))_{1:50,1:50}$ (B) The heatmap of $(\text{cov2cor}(C^{-1}))_{1:50,1:50}$

FIGURE 2. The snapshots of the first 50 columns and rows of the covariance and inverse covariance matrix showcasing strong correlation structure and lack of strong sparsity (large numbers of empty squares)

still exhibiting better performance than that of sample mean. Figures 3a and 3b showcase that the size of the labeled samples does not impact previous findings.

TABLE 1. Comparison for the Mean Squared Error(MSE), Average Length(AL) and Average Coverage(AC) of the confidence intervals based on \bar{Y} and $\hat{\theta}$ in Model 5.1

m	n	p	s_0	$\text{MSE}(\bar{Y})$	$\text{MSE}(\hat{\theta})$	$\text{AL}(\bar{Y})$	$\text{AL}(\hat{\theta})$	$\text{AC}(\bar{Y})$	$\text{AC}(\hat{\theta})$
1000	100	500	1	0.02064	0.01259	0.550	0.438	0.938	0.938
2000	200	500	1	0.00906	0.00521	0.391	0.301	0.954	0.966
5000	500	500	1	0.00411	0.00236	0.248	0.186	0.942	0.942
1000	100	500	5	0.01917	0.01632	0.553	0.518	0.956	0.968
2000	200	500	5	0.01014	0.00735	0.390	0.336	0.948	0.956
5000	500	500	5	0.00363	0.00248	0.248	0.197	0.958	0.944
1000	100	500	15	0.02226	0.02096	0.553	0.539	0.944	0.942
2000	200	500	15	0.00958	0.00845	0.393	0.369	0.956	0.950
5000	500	500	15	0.00377	0.00277	0.247	0.215	0.956	0.954


(A) The MSE of \bar{Y} and $\hat{\theta}$ as n varies(B) The MSE of S_Y^2 and $\hat{\sigma}_Y^2$ as n varies

FIGURE 3. Mean squared error of the proposed estimates (dashed) and the sample mean (left) and sample variance (right) both denoted with full lines. Colors denote different sparsity levels of the l_2 projections. Data was generated according to the Model 5.1 with $p = 500$, $m = 10n$.

TABLE 2. Comparison for the Mean Squared Error(MSE), Average Length(AL) and Average Coverage(AC) of the confidence intervals based on S_Y^2 and $\hat{\sigma}_Y^2$ in Model 5.1

m	n	p	s_0	$\text{MSE}(S_Y^2)$	$\text{MSE}(\hat{\sigma}_Y^2)$	$\text{AL}(S_Y^2)$	$\text{AL}(\hat{\sigma}_Y^2)$	$\text{AC}(S_Y^2)$	$\text{AC}(\hat{\sigma}_Y^2)$
1000	100	500	1	0.09129	0.07836	1.111	0.991	0.930	0.914
2000	200	500	1	0.03829	0.03041	0.781	0.694	0.946	0.940
5000	500	500	1	0.01615	0.01313	0.498	0.439	0.942	0.938
1000	100	500	5	0.07513	0.07176	1.113	1.058	0.938	0.934
2000	200	500	5	0.04140	0.03862	0.778	0.736	0.928	0.922
5000	500	500	5	0.01493	0.01281	0.499	0.457	0.950	0.942
1000	100	500	15	0.07763	0.07553	1.112	1.106	0.948	0.932
2000	200	500	15	0.03739	0.03656	0.792	0.769	0.952	0.952
5000	500	500	15	0.01587	0.01466	0.496	0.476	0.940	0.942

Model 5.2. Assume X_1 is non-Gaussian and the linear model is not correct. We consider

$$X_{ij} \sim^{\text{iid}} \text{Pois}(1), \quad \forall (i, j) \in \{1, 2, \dots, m+n\} \times \{1, 2, \dots, p-1\}$$

and

$$Y_i = 0.1 * (X_i^\top \beta^0)^2 + X_i^\top \beta^0 + \delta_i, \quad i = 1, 2, \dots, n,$$

where $\beta^0 = (\frac{1}{\sqrt{s_0}} \mathbf{1}_{s_0}^\top, 0^\top)^\top$, $\delta_i \sim^{\text{iid}} \mathcal{N}(0, 1)$. Fix $K = 2$, set $\hat{\beta}^{(-k)}$ as the square-root Lasso estimator with tuning parameter $\lambda = \sqrt{\log p / (n - N)}$. Let $\hat{\theta}$ and $\hat{\sigma}_Y^2$ be the semi-supervised estimators for θ and σ_Y^2 defined as in (4) and (7), respectively. Then, we propose the confidence intervals for θ and σ_Y^2 as in (8) and (19).

We repeat the simulation for 500 times. Table 3 shows the mean squared error (MSE), the average length (AL) and the average coverage (AC) of confidence intervals for θ estimated by \bar{Y} and $\hat{\theta}$. Besides, Table 4 shows the mean squared error (MSE), the average length (AL) and the average coverage (AC) of confidence intervals for σ_Y^2 estimated by S_Y^2 and $\hat{\sigma}_Y^2$. Figure 4a and 4b show the MSE of \bar{Y} , $\hat{\theta}$, S_Y^2 and $\hat{\sigma}_Y^2$ as n varies for several choices of the sparsity s_0 .

It is worth noting that the Lasso estimate of the actual parameter β^0 is not consistent in this quadratic model; Lasso was fitted assuming a linear model is valid. Nevertheless, our estimate of the mean of the response is still universally better than that of sample mean; smaller MSE, shorter length as well as close to nominal coverage probability. Estimates of the second moment and therefore variance, deteriorate more with significant departures from the model; however if the sparsity is strongly present in the model, we still see a reduction in variability compared to the sample variance estimate.

TABLE 3. Comparison for the Mean Squared Error(MSE), Average Length(AL) and Average Coverage(AC) of the confidence intervals based on S_Y^2 and $\hat{\sigma}_Y^2$ in Model 5.2

m	n	p	s_0	MSE(\bar{Y})	MSE($\hat{\theta}$)	AL(\bar{Y})	AL($\hat{\theta}$)	AC(\bar{Y})	AC($\hat{\theta}$)
1000	100	500	1	0.02406	0.01475	0.562	0.452	0.916	0.932
2000	200	500	1	0.01151	0.00620	0.396	0.303	0.930	0.942
5000	500	500	1	0.00407	0.00217	0.250	0.187	0.958	0.956
1000	100	500	5	0.02141	0.02055	0.564	0.547	0.936	0.932
2000	200	500	5	0.01122	0.00834	0.401	0.353	0.930	0.946
5000	500	500	5	0.00410	0.00237	0.254	0.198	0.954	0.960
1000	100	500	15	0.02057	0.02104	0.575	0.580	0.952	0.944
2000	200	500	15	0.01139	0.01078	0.407	0.397	0.942	0.946
5000	500	500	15	0.00443	0.00346	0.258	0.229	0.946	0.950

TABLE 4. Comparison for the Mean Squared Error(MSE), Average Length(AL) and Average Coverage(AC) of the confidence intervals based on S_Y^2 and $\hat{\sigma}_Y^2$ in Model 5.2

m	n	p	s_0	MSE(S_Y^2)	MSE($\hat{\sigma}_Y^2$)	AL(S_Y^2)	AL($\hat{\sigma}_Y^2$)	AC(S_Y^2)	AC($\hat{\sigma}_Y^2$)
1000	100	500	1	0.09017	0.07865	1.154	1.090	0.928	0.932
2000	200	500	1	0.04988	0.03807	0.808	0.750	0.918	0.944
5000	500	500	1	0.01906	0.01365	0.508	0.455	0.932	0.942
1000	100	500	5	0.09312	0.10042	1.163	1.177	0.930	0.932
2000	200	500	5	0.03928	0.03866	0.827	0.816	0.954	0.958
5000	500	500	5	0.01794	0.01518	0.523	0.484	0.948	0.948
1000	100	500	15	0.09086	0.10350	1.205	1.218	0.942	0.930
2000	200	500	15	0.04976	0.05225	0.850	0.859	0.940	0.944
5000	500	500	15	0.02017	0.01981	0.539	0.532	0.956	0.948


(A) The MSE of \bar{Y} and $\hat{\theta}$ as n varies(B) The MSE of S_Y^2 and $\hat{\sigma}_Y^2$ as n varies

FIGURE 4. Mean squared error of the proposed estimates (dashed) and the sample mean (left) and sample variance (right) both denoted with full lines. Colors denote different sparsity levels of the l_2 projections. Data was generated according to the Model 5.2 with $p = 500$, $m = 10n$.

Model 5.3. We compare the behavior of our semi-supervised estimators under different choices of K 's when the linear model is correct. We generate $Z_{ij} \sim^{\text{iid}} \text{Exp}(1) - 1$ for all $(i, j) \in \{1, 2, \dots, m+n\} \times \{1, 2, \dots, p-1\}$. Let $X_i = C_1^{\frac{1}{2}} Z_i + 1_{p-1}$, where the covariance matrix C_1 is defined as in Model 5.1. Set


$$Y_i = \sum_{i=1}^{s_0} X_i / \sqrt{s_0} + \mathcal{N}(0, 1).$$

Fix $n = 48$, $p = 50$, $m = 96$, $s_0 \in \{1, 3, 5\}$ and let K varies from $\{2, 3, 4, 6, 8, 12, 16, 24, 48\}$. Set $\hat{\beta}^{(-k)}$ as the Lasso estimator with 10-fold cross-validation. Let $\hat{\theta}$ and $\hat{\sigma}_Y^2$ be the semi-supervised estimators for θ and σ_Y^2 defined as in (4) and (7), respectively. We repeat the simulation for 500 times, Figure 5a and 5b show the MSE's of $\hat{\theta}$ and $\hat{\sigma}_Y^2$ as K varies for several choices of s_0 . We observe that with larger K , MSEs do not increase and in fact tend to stabilize with larger values of K .

Model 5.4. Now we compare the behavior of the semi-supervised estimators under different choices of K 's when the linear model is not correct and the sparsity s_0 is relatively large. Generate $Z_{ij} \sim^{\text{iid}} \text{Pois}(1) - 1$ for all $(i, j) \in \{1, 2, \dots, m+n\} \times \{1, 2, \dots, p-1\}$. Let $X_i = C_1^{\frac{1}{2}} Z_i + 1_{p-1}$, where the covariance matrix C_1 is defined as in Model 5.1. Set


$$Y_i = 0.1 \left(\sum_{i=1}^{s_0} X_i / \sqrt{s_0} \right)^2 + \sum_{i=1}^{s_0} X_i / \sqrt{s_0} + \mathcal{N}(0, 1).$$

Fix $n = 48$, $m = 96$, $s_0 = \{5, 10, 15\}$ and let K varies from $\{2, 3, 4, 6, 8, 12, 16, 24, 48\}$. Set $\hat{\beta}^{(-k)}$ as the square-root Lasso estimator with $\lambda = \sqrt{\log p / (n - N)}$. Let $\hat{\theta}$ and $\hat{\sigma}_Y^2$ be the semi-supervised estimators for $\mathbb{E}Y_1$ and $\text{Var}Y_1$ defined as in (4) and (7), respectively. We repeat the simulation for 500 times, Figure 6a and 6b show the MSE's of $\hat{\theta}$ and $\hat{\sigma}_Y^2$ as



(A) The MSE of $\hat{\theta}$ as K varies, Model 5.3 (B) The MSE of $\hat{\sigma}_Y^2$ as K varies, Model 5.3

FIGURE 5. Stability of the proposed estimates as number of folds, K , varies leading to lastly the leave-one-out updates. Data used $n = 48$, $m = 96$. Different colors correspond to different sparsity levels of the l_2 projection. The left plot corresponds to the mean whereas the right to the variance.



(A) The MSE of $\hat{\theta}$ as K varies, Model 5.4 (B) The MSE of $\hat{\sigma}_Y^2$ as K varies, Model 5.4

FIGURE 6. Stability of the proposed estimates as number of folds, K , varies leading to lastly a leave-one-out updates. Data used $n = 48$, $m = 96$. Different colors correspond to different sparsity levels of the l_2 projection. The left plot corresponds to the mean whereas the right to the variance.

K varies for several choices of s_0 . Therein, it seems clear that both estimates have nearly decreasing MSE's; note that with $K = 48$, we obtain leave-one-out estimates.

5.2. Real data. We use our semi-supervised procedure to estimate the average treatment effect and the Treatment Effect Size of a genotype mutation on Human Immunodeficiency Virus (HIV) drug resistance.

We consider the dataset derived from Baxter et al. (2006), the dataset is available on the Stanford University HIV Drug Resistance Database (Rhee et al., 2003) (<https://hivdb.stanford.edu/pages/genopheno.dataset.html>). The outcome variable Y is the drug resistance, here we consider the 3TC (lamivudine) nucleoside/nucleotide RT inhibitor. The treatment Z_k is an identity variable denotes whether there is a mutation on the T -th position of the HIV reverse transcriptase (RT). Each of the explanatory

variable X_j ($j \in \{1, 2, \dots, 240\} \setminus \{T\}$) is also an identity variable which denotes whether there is a mutation on the j -th position. Here, we only consider the subtype B sequence. Besides, redundant viruses obtained from the same individual were excluded. We are assuming that the the outcome variable is independent of whether it is observable or not.

From the dataset, we have $n = 518$ pairs of supervised data ($\{X_{i,j}\}_{j \neq T}, Z_{i,T}, Y_i$) ($i \in \{1, 2, \dots, n\}$) and $m = 2428$ pairs of additional unlabeled covariates ($\{X_{i,j}\}_{j \neq T}, Z_{i,T}$) ($i \in \{n+1, n+2, \dots, m+n\}$). Fix $T \in \{1, 2, \dots, 240\}$. We only consider the covariates that the corresponding mutations occur for at least 10 times in the supervised data for each group (treatment and control).

We apply the semi-supervised procedure as in Section 4. Split the supervised dataset into $K = 3$ folds, let $\widehat{\beta}_0^{(-k)}$ and $\widehat{\beta}_1^{(-k)}$ be the Lasso estimators for the slopes, the tuning parameters are chosen by 10-fold cross-validation. As for the choice of $\widehat{e}^{(-k)}$, we use 1) Logistic Regression (LR) with Lasso penalty and 10-fold cross-validation; 2) Support Vector Machine (SVM) with linear kernel; 3) Random Forest (RF). In order to improve the stability of the estimator, we trim each $\widehat{e}^{(-k)}(X_i)$ to $[0.01, 0.99]$. Let $\widehat{\delta}_s$ and \widehat{d}_s denote the sample average treatment effect and the sample Treatment Effect Size, respectively. Let $\widehat{\delta}_{\text{LR}}, \widehat{\delta}_{\text{SVM}}$ and $\widehat{\delta}_{\text{RF}}$ denote the semi-supervised estimators for the ATE. Let $\widehat{d}_{\text{LR}}, \widehat{d}_{\text{SVM}}$ and \widehat{d}_{RF} denote the semi-supervised estimators for the TES. We compare our results with Chernozhukov's supervised estimator. [Chernozhukov et al. \(2018\)](#) In order to reduce the variation derived from the randomness of sample splitting, we repeat the sample splitting procedure for $t_{\max} = 5$ times.

Figure 7 and 8 show the confidence intervals for δ and d on several positions based on the different estimators. There are several relatively long confidence intervals in Figure 7 and 8, such as $\widehat{\delta}_{\text{RF}}$ on position 98, $\widehat{\delta}_{\text{LR}}$ and $\widehat{\delta}_{\text{SVM}}$ on position 151. This is because the procedure is not stable enough that some of the $\widehat{e}^{(-k)}(X_i)$'s are close to 0 or 1.

We can see that, there is a very large TES on position 184, a small TES on position 69, 75, 151 and potentially a very small TES on position 203 and 215. As stated in [Rhee et al. \(2003\)](#), nearly all drug-resistance mutations are found between positions 40-240. Moreover, the NRTI Mutation Pattern and Susceptibility can be found

<https://hivdb.stanford.edu/pages/phenoSummary/Pheno.NRTI.Simple.html>.

On position 41, 44, 74, 98 and 203, the sample estimator $\widehat{\delta}_s$ is obviously different from the supervised and semi-supervised estimator of ATE based on propensity score. This is because that the sample estimator is estimating

$$\mathbb{E}(Y_1|D_1 = 1) - \mathbb{E}(Y_1|D_1 = 0) = \beta_1^{*\top} \mathbb{E}(X_1|D_1 = 1) - \beta_0^{*\top} \mathbb{E}(X_1|D_1 = 0),$$

which is different from $\delta = \mathbb{E}\mathbb{E}(Y_1|D_1 = 1, X) - \mathbb{E}\mathbb{E}(Y_1|D_1 = 0, X) = (\beta_1^* - \beta_0^*)^\top \mathbb{E}X_1$ when $\mathbb{E}(X_1|D_1 = 1) \neq \mathbb{E}(X_1|D_1 = 0)$. This result suggests that, the mutations on position 41, 44, 74, 98 and 203 usually occur simultaneously with the mutations on other positions who potentially contribute an enhancement of the 3TC drug resistance.




FIGURE 7. CI's for the Average Treatment Effects estimated by sample estimator $\hat{\delta}_s$, supervised estimator based on Logistic Regression (LR) $\hat{\delta}_{LR}$, Support Vector Machine (SVM) $\hat{\delta}_{SVM}$ and Random Forest (RF) $\hat{\delta}_{RF}$; 3) semi-supervised estimator based on LR $\hat{\delta}_{SSLR}$, SVM $\hat{\delta}_{SSSVM}$ and RF $\hat{\delta}_{SSRF}$




FIGURE 8. CI's for the Treatment Effect Size estimated by sample estimator \hat{d}_s , supervised estimator based on Logistic Regression (LR) \hat{d}_{LR} , Support Vector Machine (SVM) \hat{d}_{SVM} and Random Forest (RF) \hat{d}_{RF} ; 3) semi-supervised estimator based on LR \hat{d}_{SSLR} , SVM \hat{d}_{SSSVM} and RF \hat{d}_{SSRF}

6. DISCUSSION

In this paper, we presented a new way of estimating the mean and variance of the outcomes of interests in the setting of semi-supervised observations. Observations are allowed to be much larger compared to the size of the labeled outcomes. The key idea is to replace the linear or quadratic forms with appropriate bias correction terms that try to utilize as much as possible the dependence of the outcomes to the covariates. Additional observations are utilized to estimate the first and second moments of the covariate distribution. We showed that this strategy leads to significant improvements in efficiency (reduction in asymptotic variance) of the naive estimates. Moreover, we identified a set of conditions in a high-dimensional setting that is weaker than those of methods depending on only labeled observations and accommodate for large uncertainties in the selection of possibly important variables. Our work builds upon similar results obtained in [Zhang et al. \(2016\)](#) in the low-dimensional setting. However, we have observed that our simple, modified estimator improves on their results by providing better control even in the low-dimensional setting. This allows greater flexibility in estimating the dependence of the outcomes and covariates and leaves many open questions; for example, can fully non-parametric methods like those of random forests, deep neural networks and such be also utilized?

REFERENCES

- Azriel, D., L. D. Brown, M. Sklar, R. Berk, A. Buja, and L. Zhao (2016). Semi-supervised linear regression. *arXiv preprint arXiv:1612.02391*.
- Barber, R. F., E. J. Candès, et al. (2015). Controlling the false discovery rate via knockoffs. *The Annals of Statistics* 43(5), 2055–2085.
- Baxter, J. D., J. M. Schapiro, C. A. Boucher, V. M. Kohlbrenner, D. B. Hall, J. R. Scherer, and D. L. Mayers (2006). Genotypic changes in human immunodeficiency virus type 1 protease associated with reduced susceptibility and virologic response to the protease inhibitor tipranavir. *Journal of virology* 80(21), 10794–10801.
- Belloni, A., V. Chernozhukov, and L. Wang (2011). Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika* 98(4), 791–806.
- Bogdan, M., E. Van Den Berg, C. Sabatti, W. Su, and E. J. Candès (2015). Slope—adaptive variable selection via convex optimization. *The annals of applied statistics* 9(3), 1103.
- Bühlmann, P., S. van de Geer, et al. (2015). High-dimensional inference in misspecified linear models. *Electronic Journal of Statistics* 9(1), 1449–1473.
- Cai, T. T. and Z. Guo (2018). Semi-supervised inference for explained variance in high-dimensional linear regression and its applications. *arXiv preprint arXiv:1806.06179*.
- Cai, T. T., Z. Guo, et al. (2017). Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *The Annals of statistics* 45(2), 615–646.
- Cannings, T. I., T. B. Berrett, and R. J. Samworth (2017). Local nearest neighbour classification with applications to semi-supervised learning. *arXiv preprint arXiv:1704.00642*.

- Chakraborty, A. and T. Cai (2017). Efficient and adaptive linear regression in semi-supervised settings. *arXiv preprint arXiv:1701.04889*.
- Chapelle, O., B. Scholkopf, and A. Zien (2009). Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks* 20(3), 542–542.
- Cheng, D., A. Ananthakrishnan, and T. Cai (2018). Efficient and robust semi-supervised estimation of average treatment effects in electronic medical records data. *arXiv preprint arXiv:1804.00195*.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, and W. Newey (2017). Double/debiased/neyman machine learning of treatment effects. *American Economic Review* 107(5), 261–65.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1), C1–C68.
- De Mol, C., E. De Vito, and L. Rosasco (2009). Elastic-net regularization in learning theory. *Journal of Complexity* 25(2), 201–230.
- Dicker, L. H. (2014). Variance estimation in high-dimensional linear models. *Biometrika* 101(2), 269–284.
- El Alaoui, A., X. Cheng, A. Ramdas, M. J. Wainwright, and M. I. Jordan (2016). Asymptotic behavior of ℓ_p -based laplacian regularization in semi-supervised learning. In *Conference on Learning Theory*, pp. 879–906.
- Fan, J., S. Guo, and N. Hao (2012). Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74(1), 37–65.
- Grandvalet, Y. and Y. Bengio (2005). Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pp. 529–536.
- Gronsbell, J. L. and T. Cai (2018). Semi-supervised approaches to efficient evaluation of model prediction performance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80(3), 579–594.
- Hebiri, M., S. Van De Geer, et al. (2011). The smooth-lasso and other $\ell_1 + \ell_2$ -penalized methods. *Electronic Journal of Statistics* 5, 1184–1226.
- Holland, P. W. (1988). Causal inference, path analysis and recursive structural equations models. *ETS Research Report Series 1988(1)*, i–50.
- Hummel, R., S. Banga, and T. P. Hettmansperger (2005). Better confidence intervals for the variance in a random sample. Technical report, Citeseer.
- Janson, L., R. F. Barber, and E. Candes (2017). Eigenprism: inference for high dimensional signal-to-noise ratios. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(4), 1037–1065.
- Kim, K. I., F. Steinke, and M. Hein (2009). Semi-supervised regression using hessian energy with an application to semi-supervised dimensionality reduction. In *Advances in Neural Information Processing Systems*, pp. 979–987.

- Kuchibhotla, A. K., L. D. Brown, A. Buja, E. I. George, and L. Zhao (2018). A model free perspective for linear regression: Uniform-in-model bounds for post selection inference. *arXiv preprint arXiv:1802.05801*.
- Künzel, S. R., J. S. Sekhon, P. J. Bickel, and B. Yu (2017). Meta-learners for estimating heterogeneous treatment effects using machine learning. *arXiv preprint arXiv:1706.03461*.
- Mai, X. and R. Couillet (2017). A random matrix analysis and improvement of semi-supervised learning for large dimensional data. *arXiv preprint arXiv:1711.03404*.
- Moscovich, A., A. Jaffe, and B. Nadler (2016). Minimax-optimal semi-supervised regression on unknown manifolds. *arXiv preprint arXiv:1611.02221*.
- Neyman, J. S. (1923). On the application of probability theory to agricultural experiments. essay on principles. section 9.(translated and edited by dm dabrowska and tp speed, statistical science (1990), 5, 465-480). *Annals of Agricultural Sciences* 10, 1–51.
- Raskutti, G., M. J. Wainwright, and B. Yu (2010). Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research* 11(Aug), 2241–2259.
- Rhee, S.-Y., M. J. Gonzales, R. Kantor, B. J. Betts, J. Ravela, and R. W. Shafer (2003). Human immunodeficiency virus reverse transcriptase and protease sequence database. *Nucleic acids research* 31(1), 298–303.
- Rinaldo, A., L. Wasserman, M. G’Sell, J. Lei, and R. Tibshirani (2016). Bootstrapping and sample splitting for high-dimensional, assumption-free inference. *arXiv preprint arXiv:1611.05401*.
- Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* 66(5), 688.
- Stucky, B. and S. Van De Geer (2017). Sharp oracle inequalities for square root regularization. *The Journal of Machine Learning Research* 18(1), 2256–2284.
- Sun, Q., W.-X. Zhou, and J. Fan (2018). Adaptive huber regression. *Journal of the American Statistical Association* (just-accepted), 1–35.
- Van De Geer, S. A., P. Bühlmann, et al. (2009). On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics* 3, 1360–1392.
- Wager, S. and S. Athey (2017). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* (just-accepted).
- Wasserman, L. and J. D. Lafferty (2008). Statistical analysis of semi-supervised regression. In *Advances in Neural Information Processing Systems*, pp. 801–808.
- Zhang, A., L. D. Brown, and T. T. Cai (2016). Semi-supervised inference: General theory and estimation of means. *arXiv preprint arXiv:1606.07268*.
- Zhou, Z.-H. and M. Li (2005). Semi-supervised regression with co-training. In *IJCAI*, Volume 5, pp. 908–913.

Zhu, X. (2005). Semi-supervised learning literature survey.

Zhu, X. and A. B. Goldberg (2006). Semi-supervised regression with order preferences. *Dept. Comput. Sci., Univ. Wisconsin-Madison, Madison, WI, USA, Tech. Rep 1578*, 10.

Zhu, X. and A. B. Goldberg (2009). Introduction to semi-supervised learning. *Synthesis lectures on artificial intelligence and machine learning* 3(1), 1–130.

APPENDIX A. PROOFS OF THE MAIN RESULTS

All results and notation are numbered and used, as in the main text unless stated otherwise. Recall, that for any $a \in \mathbb{R}^p$ and $A \in \mathbb{R}^{p \times p}$, let $\|a\|_A = a^\top A a$. Throughout the article constants $c, C > 0$, independent of n, p, m , may change value from line to line.

With a slight abuse of notation, for any function g ,

$$\mathbb{E}_{I_k^c} g = \mathbb{E}(g | (Y_i, X_i)_{i \in \{1, 2, \dots, m+n\} \setminus I_k})$$

in the proof of Theorems 1, 2, 3, 4, 5;

$$\mathbb{E}_{I_k^c} g = \mathbb{E}(g | (D_i, Y_i, X_i)_{i \in \{1, 2, \dots, m+n\} \setminus I_k})$$

in the proof of Theorems 6, 7, 8, 9.

We begin by presenting two simple results that we will be useful throughout the document.

Lemma 1 (Lemma B.1 in Chernozhukov et al. (2017)). *Let $\{X_n\}$ and $\{Y_n\}$ be sequences of random variables. If for any $c > 0$, $\mathbb{P}(|X_n| > c|Y_n) = o_P(1)$. Then, $X_n = o_P(1)$. In particular, this occurs if $\mathbb{E}(|X_n|^q | Y_n) = o_P(1)$ for any $q \geq 1$. Typical examples we used in our proofs are a) $\mathbb{E}(X_n^2 | Y_n) = o_P(1)$, b) $X_n = \sum_{i=1}^n Z_{n,i}/n$, where $(Z_{n,i})$ is a row-wise independent and identical distributed triangular array with $\mathbb{E}(|Z_{n,1}| | Y_n) = o_P(1)$.*

Lemma 2. *Let $\{X_n\}$ and $\{Y_n\}$ be sequences of random variables. If for any $c > 0$, there exists $M, N > 0$ such that for any $n > N$, $\mathbb{P}(|X_n| > M | Y_n) < c$. Then, $X_n = \mathcal{O}_P(1)$. In particular, this occurs if $\mathbb{E}X_n = 0$ and $\mathbb{E}(X_n^2 | Y_n) < C$ for some $C < \infty$.*

Proof of Lemma 2. For any $c > 0$,

$$\mathbb{P}(|X_n| > M) = \mathbb{E}(\mathbb{P}(|X_n| > M | Y_n)) < c, \quad \forall n > N$$

That is, $X_n = \mathcal{O}_P(1)$. If $\mathbb{E}X_n = 0$ and $\mathbb{E}(|X_n|^2 | Y_n) < C$, then by Chebyshev's Inequality,

$$\mathbb{P}(|X_n| \geq \sqrt{C/c} | Y_n) \leq \mathbb{P}(|X_n| \geq \sqrt{\text{Var}X_n/c}) \leq c.$$

Lemma 3. *Let $(Z_{n,i})$ be a row-wise independent and identical distributed triangular array with $\mathbb{E}Z_{n,1} = 0$ and $\mathbb{E}|Z_{n,1}|^q < C$ for $q \geq 1$ and $C < \infty$. Let $X_n = \sum_{i=1}^n Z_{n,i}/n$. Then, $X_n = o_P(1)$.*

Proof of Lemma 3. Let $Y_{n,i} = Z_{n,i}1_{|Z_{n,i}| < n}$. For any $c > 0$,

$$\mathbb{P}(|X_n| \geq c) \leq \mathbb{P}(\cup_{i=1}^n [Z_{n,i} \neq Y_{n,i}]) + \mathbb{P}(|\sum_{i=1}^n Y_{n,i}| \geq nc)$$

Let $r \in (1, q \wedge 2)$, then $\mathbb{E}|Z_{n,1}|^r \leq (\mathbb{E}|Z_{n,1}|^q)^{r/q} < C^{r/q}$. By Markov's Inequality,

$$\begin{aligned} \mathbb{P}(\cup_{i=1}^n [Z_{n,i} \neq Y_{n,i}]) &\leq n\mathbb{P}(|Z_{n,1}| \geq n) \leq n\mathbb{E}|Z_{n,1}|^q/n^q = n^{1-q}\mathbb{E}|Z_{n,1}|^q = o(1), \\ \mathbb{P}(|\sum_{i=1}^n Y_{n,i}| \geq nc) &\leq \mathbb{E}|\sum_{i=1}^n Y_{n,i}|^2/(nc)^2 = n\mathbb{E}Z_{n,1}^2 1_{[|Z_{n,1}| < n]}/(nc)^2 \\ &= \mathbb{E}|Z_{n,1}|^r |Z_{n,1}|^{2-r} 1_{[|Z_{n,1}| < n]}/(nc^2) \leq n^{1-r}\mathbb{E}|Z_{n,1}|^r/c^2 = o(1). \end{aligned}$$

Hence,

$$\mathbb{P}(|X_n| \geq c) = o_P(1).$$

That is, $X_n = o_P(1)$.

Proof of Theorem 1. Theorem 1 is a simple corollary of all of the main results as presented in Theorems 2, 3, 4, 5.

Proof of Theorem 2. This proof provides \sqrt{n} consistencies of $\hat{\theta}$ and $\hat{\sigma}_Y^2$. We first assume Condition 1 and 3 and show that $\hat{\theta} - \theta = \mathcal{O}_P(n^{-\frac{1}{2}})$. By the definition of $\hat{\theta}^{(-k)}$ as in (2), we can see that

$$(52) \quad \hat{\theta}^{(-k)} - \theta = N^{-1} \sum_{i \in I_k} (Y_i - \theta) - N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i + m^{-1} \sum_{i=n+1}^{m+n} \hat{\beta}^{(-k)\top} \tilde{V}_i.$$

By the definition of β^* , as in Lemma 1 of Zhang et al. (2016), we have

$$\mathbb{E}\varepsilon = 0, \quad \mathbb{E}\tilde{X}\varepsilon = 0, \quad \theta = \beta^{*\top} \tilde{\mu}, \quad \sigma_Y^2 = b^2 + \varepsilon^2.$$

From Condition 1 and 3, we have

$$\begin{aligned} \mathbb{E}(Y - \theta)^2 &< (\mathbb{E}|Y - \theta|^{2+c})^{\frac{2}{2+c}} < C, \\ \mathbb{E}(\beta^{*\top} \tilde{V})^2 &= \mathbb{E}(Y - \theta)^2 - \sigma_\varepsilon^2 \leq \mathbb{E}(Y - \theta)^2 < C, \\ \mathbb{E}_{I_k^c} \left((\hat{\beta}^{(-k)} - \beta^*)^\top \tilde{V} \right)^2 &= \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}} \mathbb{E}_{I_k^c} \left((C^{\frac{1}{2}} (\hat{\beta}_{-1}^{(-k)} - \beta_{-1}^*))^\top Z / \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}} \right)^2 \\ &\leq \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}} \sup_{\|a\|_2=1} \mathbb{E}(a^\top Z)^2 \\ &\leq \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}} \sup_{\|a\|_2=1} (\mathbb{E}|a^\top Z|^{2+c})^{\frac{2}{2+c}} < C. \end{aligned}$$

and hence by triangle inequality,

$$\mathbb{E}_{I_k^c} (\hat{\beta}^{(-k)\top} \tilde{V})^2 < C.$$

Then, it is easy to see

$$(53) \quad N^{-1} \sum_{i \in I_k} (Y_i - \theta) = \mathcal{O}_P(N^{-\frac{1}{2}}),$$

$$(54) \quad N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widetilde{V}_i = \mathcal{O}_P(N^{-\frac{1}{2}}),$$

$$(55) \quad m^{-1} \sum_{i=n+1}^{m+n} \widehat{\beta}^{(-k)\top} \widetilde{V}_i = \mathcal{O}_P(m^{-\frac{1}{2}}).$$

Therefore,

$$\widehat{\theta}^{(k)} - \theta = \mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(m^{-\frac{1}{2}}) = \mathcal{O}_P(N^{-\frac{1}{2}}).$$

When $K < \infty$,

$$(56) \quad \widehat{\theta} = \theta + \mathcal{O}_P(n^{-\frac{1}{2}}).$$

Now we assume Condition 2 and 3 and show that $\widehat{\sigma}_Y^2 - \sigma_Y^2 = \mathcal{O}_P(n^{-\frac{1}{2}})$. By the definition of $\widehat{\sigma}_Y^{2(k)}$ as in (5), we can see that

$$\widehat{\sigma}_Y^{2(k)} = N^{-1} \sum_{i \in I_k} (Y_i - \widehat{\theta})^2 + N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \left(\widehat{C} - \widehat{V}_i^\top \widehat{V}_i \right) \widehat{\beta}^{(-k)},$$

By (53)-(56) and the fact that $a^2 = b^2 + (a-b)^2 + 2(a-b)b$,

$$(57) \quad \begin{aligned} N^{-1} \sum_{i \in I_k} (Y_i - \widehat{\theta})^2 &= N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 + (\widehat{\theta} - \theta)^2 - 2(\widehat{\theta} - \theta) N^{-1} \sum_{i \in I_k} (Y_i - \theta) \\ &= N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 + \mathcal{O}_P(n^{-1}) + \mathcal{O}_P(n^{-\frac{1}{2}}) \cdot \mathcal{O}_P(N^{-\frac{1}{2}}), \\ &= N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 + \mathcal{O}_P(N^{-1}), \\ \widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)} &= m^{-1} \sum_{i=n+1}^{m+n} \left(\widehat{\beta}^{(-k)\top} (\widetilde{X}_i - \widetilde{\mu}) \right)^2 - \left(\widehat{\beta}^{(-k)\top} (\widehat{\mu} - \widetilde{\mu}) \right)^2 \\ &= m^{-1} \sum_{i=n+1}^{m+n} \left(\widehat{\beta}^{(-k)\top} \widetilde{V}_i \right)^2 + \mathcal{O}_P(m^{-1}), \\ N^{-1} \sum_{i \in I_k} \left(\widehat{\beta}^{(-k)\top} \widehat{V}_i \right)^2 &= N^{-1} \sum_{i \in I_k} \left(\widehat{\beta}^{(-k)\top} \widetilde{V}_i \right)^2 + \left(\widehat{\beta}^{(-k)\top} (\widehat{\mu} - \widetilde{\mu}) \right)^2 \\ &\quad - 2 \widehat{\beta}^{(-k)\top} (\widehat{\mu} - \widetilde{\mu}) N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widetilde{V}_i \\ &= N^{-1} \sum_{i \in I_k} \left(\widehat{\beta}^{(-k)\top} \widetilde{V}_i \right)^2 + \mathcal{O}_P(m^{-1} + N^{-\frac{1}{2}} m^{-\frac{1}{2}}). \end{aligned}$$

Hence,

$$(58) \quad \widehat{\sigma}_Y^{2^{(k)}} = N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 - N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(N^{-1}).$$

Similarly as before, utilizing Condition 2, we observe

$$\mathbb{E}(Y - \theta)^4 < C, \quad \mathbb{E}(\beta^{*\top} \widetilde{V})^4 < C, \quad \mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V})^4 < C$$

and with that

$$\mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} \widetilde{V})^4 < C, \quad \text{Var}_{I_k^c}(\widehat{\beta}^{(-k)\top} \widetilde{V})^2 \leq \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} \widetilde{V})^4 < C.$$

Therefore, by Lemma 2,

$$\begin{aligned} N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 &= \sigma_Y^2 + \mathcal{O}_P(N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 &= \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(N^{-\frac{1}{2}}), \\ m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 &= \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(m^{-\frac{1}{2}}). \end{aligned}$$

Lastly, we obtain

$$\widehat{\sigma}_Y^{2^{(k)}} - \sigma_Y^2 = \mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(m^{-\frac{1}{2}}) + \mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(N^{-1}) = \mathcal{O}_P(N^{-\frac{1}{2}}).$$

Now the proof is finalized by noticing that for finite K , the rate above is inherited for the averaged estimator

$$\widehat{\sigma}_Y^2 = \sigma_Y^2 + \mathcal{O}_P(n^{-\frac{1}{2}}).$$

Proof of Theorem 3. This proof provides an asymptotic normal result for $\sqrt{n}(\widehat{\theta} - \theta)$ by assuming Condition 1 and 4. Recall (52),

$$\widehat{\theta}^{(k)} - \theta = N^{-1} \sum_{i \in I_k} (Y_i - \theta) - N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widetilde{V}_i + m^{-1} \sum_{i=n+1}^{m+n} \widehat{\beta}^{(-k)\top} \widetilde{V}_i.$$

Since

$$\begin{aligned} \text{Var}_{I_k^c} \left(N^{-\frac{1}{2}} \sum_{i \in I_k} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i \right) &= \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 = o_P(1), \\ \text{Var}_{I_k^c} \left(m^{-\frac{1}{2}} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i \right) &= \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 = o_P(1), \end{aligned}$$

by Lemma 1,

$$\begin{aligned} N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i &= o_P(N^{-\frac{1}{2}}), \\ m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i &= o_P(m^{-\frac{1}{2}}). \end{aligned}$$

Therefore,

$$\begin{aligned} \widehat{\theta}^{(k)} - \theta &= N^{-1} \sum_{i \in I_k} (Y_i - \theta) - N^{-1} \sum_{i \in I_k} \beta^{*\top} \widetilde{V}_i + m^{-1} \sum_{i=n+1}^{m+n} \beta^{*\top} \widetilde{V}_i + o_P(N^{-\frac{1}{2}}) + o_P(m^{-\frac{1}{2}}) \\ &= N^{-1} \sum_{i \in I_k} \varepsilon_i + m^{-1} \sum_{i=n+1}^{m+n} \beta^{*\top} \widetilde{V}_i + o_P(N^{-\frac{1}{2}}). \end{aligned}$$

By Condition 1,

$$\mathbb{E}|\varepsilon|^{2+c} < C, \quad \mathbb{E}|\beta^{*\top} \widetilde{V}|^{2+c} < C.$$

With a slight abuse of notation, assume that

$$\sigma_\varepsilon^2 = \lim_{n \rightarrow \infty} \mathbb{E}\varepsilon^2, \quad \tau b^2 = \lim_{n \rightarrow \infty} \frac{n}{m} \mathbb{E}(\beta^{*\top} \widetilde{V})^2$$

both exists.

a) When $\sigma_\varepsilon^2 > 0$ and $\tau b^2 > 0$,

$$\frac{\mathbb{E}|\varepsilon|^{2+c}}{(\mathbb{E}\varepsilon^2)^{1+c/2}} < C, \quad \frac{\mathbb{E}|\sqrt{n/m} \beta^{*\top} \widetilde{V}|^{2+c}}{(n/m \mathbb{E}(\beta^{*\top} \widetilde{V})^2)^{1+c/2}} < C,$$

i.e. the Lyapounov condition holds. By Lindeberg-Feller Central Limit Theorem,

$$n^{-\frac{1}{2}} \sum_{i=1}^n \varepsilon_i \xrightarrow{d} \mathcal{N}(0, \sigma_\varepsilon^2), \quad n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} \beta^{*\top} \widetilde{V}_i \xrightarrow{d} \mathcal{N}(0, \tau b^2).$$

By Slutsky's Theorem and multivariate delta method,

$$(59) \quad \sqrt{n}(\widehat{\theta} - \theta) = n^{-\frac{1}{2}} \sum_{i=1}^n \varepsilon_i + n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} \beta^{*\top} \widetilde{V}_i + o_P(1) \xrightarrow{d} \mathcal{N}(0, \sigma_\varepsilon^2 + \tau b^2).$$

b) When $\sigma_\varepsilon^2 = 0$, recall the assumption that $\sigma_\varepsilon^2 + \tau b^2 > 0$, we have $\tau b^2 > 0$. In this case, by Lemma 1 and Lindeberg-Feller Central Limit Theorem,

$$n^{-\frac{1}{2}} \sum_{i=1}^n \varepsilon_i = o_P(1), \quad n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} \beta^{*\top} \widetilde{V}_i \xrightarrow{d} \mathcal{N}(0, \tau b^2).$$

By Slutsky's Theorem, (59) holds.

c) When $\tau b^2 = 0$, similarly as in b), (59) holds.

Proof of Theorem 4. This proof provides consistency results for $\widehat{\sigma}_\varepsilon^2$ and \widehat{b}^2 by assuming Conditions 1 and 4. We first show that $\widehat{\sigma}_\varepsilon^2 = \sigma_\varepsilon + o_P(1)$. By the definition of $\widehat{\sigma}_\varepsilon^{2(k)}$ as in

(9),

$$\widehat{\sigma}_\varepsilon^{2^{(k)}} = N^{-1} \sum_{i \in I_k} (Y_i - \widehat{\theta} - \widehat{\beta}^{(-k)\top} \widehat{V}_i)^2.$$

By (53)-(56) and the fact that $a^2 = b^2 + (a - b)^2 + 2(a - b)b$,

$$\begin{aligned} \widehat{\sigma}_\varepsilon^{2^{(k)}} &= N^{-1} \sum_{i \in I_k} (Y_i - \theta - \widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + (\widehat{\theta} - \theta - \widehat{\beta}^{(-k)\top} (\widehat{\mu} - \widetilde{\mu}))^2 \\ &\quad - 2(\widehat{\theta} - \theta - \widehat{\beta}^{(-k)\top} (\widehat{\mu} - \widetilde{\mu})) N^{-1} \sum_{i \in I_k} (Y_i - \theta - \widehat{\beta}^{(-k)\top} \widetilde{V}_i) \\ &= N^{-1} \sum_{i \in I_k} (Y_i - \theta - \widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + (\mathcal{O}_P(n^{-\frac{1}{2}}) + \mathcal{O}_P(m^{-\frac{1}{2}}))^2 \\ &\quad + (\mathcal{O}_P(n^{-\frac{1}{2}}) + \mathcal{O}_P(m^{-\frac{1}{2}}))(\mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(N^{-\frac{1}{2}})) \\ &= N^{-1} \sum_{i \in I_k} (Y_i - \theta - \widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(N^{-1}) \end{aligned}$$

Remember that $\mathbb{E}_{I_k^c} g = \mathbb{E}(g | (Y_i, X_i)_{i \in \{1, 2, \dots, m+n\} \setminus I_k})$. By Condition 1,

$$\mathbb{E}|\varepsilon|^{2+c} < C, \quad \mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V})^2 = o_P(1)$$

and hence

$$\mathbb{E}_{I_k^c} |\varepsilon (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}| = o_P(1).$$

By Lemma 3

$$N^{-1} \sum_{i \in I_k} \varepsilon_i^2 = \sigma_\varepsilon^2 + o_P(1)$$

and by Lemma 1,

$$N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 = o_P(1), \quad N^{-1} \sum_{i \in I_k} \varepsilon_i (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i = o_P(1).$$

By the fact that $a^2 = b^2 + (a - b)^2 + 2(a - b)b$,

$$\begin{aligned} \widehat{\sigma}_\varepsilon^{2^{(k)}} &= N^{-1} \sum_{i \in I_k} (Y_i - \theta - \widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(N^{-1}) \\ &= N^{-1} \sum_{i \in I_k} \varepsilon_i^2 + N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 \\ &\quad - 2N^{-1} \sum_{i \in I_k} \varepsilon_i (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i + \mathcal{O}_P(N^{-1}) \\ &= \sigma_\varepsilon^2 + o_P(1) + o_P(1) + o_P(1) + \mathcal{O}_P(N^{-1}) = \sigma_\varepsilon^2 + o_P(1). \end{aligned}$$

Now we show that $\widehat{b}^2 = b^2 + o_P(1)$. By the definition of $\widehat{b}^{2^{(k)}}$ as in (10),

$$\widehat{b}^{2^{(k)}} = \widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)} + \frac{2}{N} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widehat{V}_i (Y_i - \widehat{\theta} - \widehat{\beta}^{(-k)\top} \widehat{V}_i).$$

By Conditions 1 and 4,

$$\mathbb{E}_{I_k^c} |\hat{\beta}^{(-k)\top} \tilde{V}|^{2+c} < C, \quad \mathbb{E}_{I_k^c} |\hat{\beta}^{(-k)\top} \tilde{V}(Y - \theta)|^{1+c/2} < C.$$

By Lemma 3,

$$\begin{aligned} m^{-1} \sum_{i=n+1}^{m+n} (\hat{\beta}^{(-k)\top} \tilde{V}_i)^2 &= \hat{\beta}^{(-k)\top} \tilde{C} \hat{\beta}^{(-k)} + o_P(1), \\ N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i (Y_i - \theta) &= \hat{\beta}^{(-k)\top} \tilde{C} \beta^* + o_P(1) = b^2 + (\hat{\beta}^{(-k)} - \beta^*)^\top \tilde{C} \beta^* + o_P(1) \\ &= b^2 + O(b \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}) + o_P(1) = b^2 + o_P(1). \end{aligned}$$

By (57) and Condition 4,

$$\begin{aligned} \hat{\beta}^{(-k)\top} \tilde{C} \hat{\beta}^{(-k)} &= m^{-1} \sum_{i=n+1}^{m+n} (\hat{\beta}^{(-k)\top} \tilde{V}_i)^2 + \mathcal{O}_P(m^{-1}) \\ &= \hat{\beta}^{(-k)\top} \tilde{C} \hat{\beta}^{(-k)} + o_P(1) + \mathcal{O}_P(m^{-1}) \\ &= \beta^{*\top} \tilde{C} \beta^* + \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}^2 + 2(\hat{\beta}^{(-k)} - \beta^*)^\top \tilde{C} \beta^* + o_P(1) + \mathcal{O}_P(m^{-1}) \\ &= \beta^{*\top} \tilde{C} \beta^* + \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}^2 + \mathcal{O}(b \|\hat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}) + o_P(1) + \mathcal{O}_P(m^{-1}) \\ &= \beta^{*\top} \tilde{C} \beta^* + o_P(1). \end{aligned}$$

Similarly,

$$N^{-1} \sum_{i \in I_k} (\hat{\beta}^{(-k)\top} \hat{V}_i)^2 = \beta^{*\top} \tilde{C} \beta^* + o_P(1).$$

By (53)-(56) and the fact that $ac = bd + (a - b)(c - d) + b(c - d) + (a - b)d$,

$$\begin{aligned} N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \hat{V}_i (Y_i - \hat{\theta}) &= N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i (Y_i - \theta) + \hat{\beta}^{(-k)\top} (\hat{\mu} - \tilde{\mu})(\hat{\theta} - \theta) \\ &\quad - \hat{\beta}^{(-k)\top} (\hat{\mu} - \tilde{\mu}) N^{-1} \sum_{i \in I_k} (Y_i - \theta) - (\hat{\theta} - \theta) N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i \\ &= N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i (Y_i - \theta) + \mathcal{O}_P(m^{-\frac{1}{2}}) \mathcal{O}_P(n^{-\frac{1}{2}}) \\ &\quad + \mathcal{O}_P(m^{-\frac{1}{2}}) \mathcal{O}_P(N^{-\frac{1}{2}}) + \mathcal{O}_P(n^{-\frac{1}{2}}) \mathcal{O}_P(N^{-\frac{1}{2}}) \\ &= N^{-1} \sum_{i \in I_k} \hat{\beta}^{(-k)\top} \tilde{V}_i (Y_i - \theta) + \mathcal{O}_P(N^{-1}) \\ &= b^2 + o_P(1) + \mathcal{O}_P(N^{-1}) = b^2 + o_P(1). \end{aligned}$$

Hence,

$$\hat{b}^{2(k)} = b^2 + o_P(1) + 2(b^2 + o_P(1)) - 2(b^2 + o_P(1)) = b^2 + o_P(1).$$

When $K < \infty$,

$$\widehat{\sigma}_\varepsilon^2 = \sigma_\varepsilon^2 + o_P(1), \quad \widehat{b}^2 = b^2 + o_P(1).$$

Now assume Conditions 2 and 4, we provide consistency rate results for $\widehat{\sigma}_Y^2$ and \widehat{b}^2 . The proof strategy is similar as above. We first consider $\widehat{\sigma}_Y^2$, by Conditions 2 and 4,

$$\mathbb{E}|\varepsilon|^{4+c} < C, \quad \mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V})^4 = o_P(1)$$

and hence

$$\mathbb{E}_{I_k^c}(\varepsilon(\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V})^2 = o_P(1).$$

By Lemma 2,

$$\begin{aligned} N^{-1} \sum_{i \in I_k} \varepsilon_i^2 &= \sigma_\varepsilon^2 + \mathcal{O}_P(N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 &= \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} \varepsilon_i (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i &= \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}} N^{-\frac{1}{2}}). \end{aligned}$$

Hence,

$$\begin{aligned} \widehat{\sigma}_\varepsilon^{2(k)} &= N^{-1} \sum_{i \in I_k} \varepsilon_i^2 + N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 \\ &\quad - 2N^{-1} \sum_{i \in I_k} \varepsilon_i (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i + \mathcal{O}_P(N^{-1}) \\ &= \sigma_\varepsilon^2 + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + N^{-\frac{1}{2}}). \end{aligned}$$

When $K < \infty$,

$$\widehat{\sigma}_Y^2 = \sigma_Y^2 + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + n^{-\frac{1}{2}}).$$

Now we show the consistency result for \widehat{b}^2 . By Conditions 2 and 4,

$$\mathbb{E}_{I_k^c} |\widehat{\beta}^{(-k)\top} \widetilde{V}|^{4+c} < C, \quad \mathbb{E}_{I_k^c} |\widehat{\beta}^{(-k)\top} \widetilde{V} (Y - \theta)|^{2+c/2} < C.$$

By Lemma 2,

$$\begin{aligned} m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 &= \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(m^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widetilde{V}_i (Y_i - \theta) &= \widehat{\beta}^{(-k)\top} \widetilde{C} \beta^* + \mathcal{O}_P(N^{-\frac{1}{2}}). \end{aligned}$$

Hence,

$$\begin{aligned}\widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)} &= m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(m^{-1}) = \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(m^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widehat{V}_i)^2 &= N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(N^{-1}) = \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widehat{V}_i (Y_i - \widehat{\theta}) &= N^{-1} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widetilde{V}_i (Y_i - \theta) + \mathcal{O}_P(N^{-1}) = \widehat{\beta}^{(-k)\top} \widetilde{C} \beta^* + \mathcal{O}_P(N^{-\frac{1}{2}}).\end{aligned}$$

Therefore,

$$\begin{aligned}\widehat{b}^{2(k)} &= \widehat{\beta}^{(-k)\top} \widehat{C} \widehat{\beta}^{(-k)} + \frac{2}{N} \sum_{i \in I_k} \widehat{\beta}^{(-k)\top} \widehat{V}_i (Y_i - \widehat{\theta} - \widehat{\beta}^{(-k)\top} \widehat{V}_i) \\ &= \widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(m^{-\frac{1}{2}}) + 2(\widehat{\beta}^{(-k)\top} \widetilde{C} \beta^* + \mathcal{O}_P(N^{-\frac{1}{2}})) \\ &\quad - 2(\widehat{\beta}^{(-k)\top} \widetilde{C} \widehat{\beta}^{(-k)} + \mathcal{O}_P(N^{-\frac{1}{2}})) \\ &= b^2 - \|\beta^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + \mathcal{O}_P(N^{-\frac{1}{2}}) = b^2 + \mathcal{O}_P(\|\beta^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + N^{-\frac{1}{2}}).\end{aligned}$$

When $K < \infty$,

$$\widehat{b}^2 = b^2 + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + n^{-\frac{1}{2}}).$$

Proof of Theorem 5. This proof provides an asymptotic normal result for $\sqrt{n}(\widehat{\sigma}_Y^2 - \sigma_Y^2)$ and a consistent estimate for the asymptotic variance. By Conditions 2 and 4,

$$\begin{aligned}\mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V})^4 &= \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^4 \mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V} / \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}})^4 \\ &\leq \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}} \sup_{\|a\|_2=1} \mathbb{E}(a^\top Z)^4 = \mathcal{O}(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^4).\end{aligned}$$

Therefore by Lemma 2,

$$\begin{aligned}N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 &= \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}}^2 N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i \beta^{*\top} \widetilde{V}_i &= (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{C} \beta^* + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widetilde{C}} N^{-\frac{1}{2}})\end{aligned}$$

Recall (58),

$$\widehat{\sigma}_Y^{2(k)} = N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 - N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 + \mathcal{O}_P(N^{-1}).$$

By the fact that $a^2 = b^2 + (a - b)^2 + 2(a - b)b$,

$$\begin{aligned}
N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 &= N^{-1} \sum_{i \in I_k} (\beta^{*\top} \widetilde{V}_i)^2 + N^{-1} \sum_{i \in I_k} ((\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i)^2 \\
&\quad + 2N^{-1} \sum_{i \in I_k} (\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{V}_i \beta^{*\top} \widetilde{V}_i \\
&= N^{-1} \sum_{i \in I_k} (\beta^{*\top} \widetilde{V}_i)^2 + \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widehat{C}}^2 + 2(\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{C} \beta^* \\
&\quad + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widehat{C}} N^{-\frac{1}{2}})
\end{aligned}$$

Similarly,

$$\begin{aligned}
m^{-1} \sum_{i=n+1}^{m+n} (\widehat{\beta}^{(-k)\top} \widetilde{V}_i)^2 &= m^{-1} \sum_{i=n+1}^{m+n} (\beta^{*\top} \widetilde{V}_i)^2 + \|\widehat{\beta}^{(-k)} - \beta^*\|_{\widehat{C}}^2 + 2(\widehat{\beta}^{(-k)} - \beta^*)^\top \widetilde{C} \beta^* \\
&\quad + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widehat{C}} m^{-\frac{1}{2}})
\end{aligned}$$

Hence,

$$\begin{aligned}
\widehat{\sigma}_Y^{2(k)} &= N^{-1} \sum_{i \in I_k} (Y_i - \theta)^2 + m^{-1} \sum_{i=n+1}^{m+n} (\beta^{*\top} \widetilde{V}_i)^2 - N^{-1} \sum_{i \in I_k} (\beta^{*\top} \widetilde{V}_i)^2 \\
&\quad + \mathcal{O}_P(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\widehat{C}} N^{-\frac{1}{2}}).
\end{aligned}$$

When $K < \infty$, by the independency between $(Y_i, X_i)_{i=1}^n$ and $\{X_i\}_{i=n+1}^{m+n}$,

$$\begin{aligned}
\sqrt{n}(\widehat{\sigma}_Y^2 - \sigma_Y^2) &= n^{-\frac{1}{2}} \sum_{i \in I_k} (\varepsilon_i^2 + 2\varepsilon_i \beta^{*\top} \widetilde{V}_i - \sigma_\varepsilon^2) + n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} ((\beta^{*\top} \widetilde{V}_i)^2 - b^2) + o_P(1) \\
&\xrightarrow{d} \mathcal{N}\left(0, \text{Var}(\varepsilon^2 + 2\beta^{*\top} \widetilde{V}\varepsilon) + \tau \text{Var}(\beta^{*\top} \widetilde{V})^2\right)
\end{aligned}$$

Now we prove the consistency of $\widehat{\sigma}_\nu^2 + \frac{n}{m} \widehat{\sigma}_\xi^2$. Recall (14),

$$\xi_i^{(k)} = \widehat{\beta}^{(-k)\top} (\widehat{V}_i \widehat{V}_i^\top - \widehat{C}) \widehat{\beta}^{(-k)}.$$

By Conditions 2 and 4,

$$\mathbb{E}_{I_k^c} |\widehat{\beta}^{(-k)\top} (\widehat{V} \widehat{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)}|^{2+c} < C.$$

By Lemma 1,

$$N^{-1} \sum_{i \in I_k} \xi_i^{(k)2} = \mathbb{E}_{I_k^c} (\widehat{\beta}^{(-k)\top} (\widehat{V} \widehat{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)})^2 + o_P(1).$$

Let $A = \tilde{V}\tilde{V}^\top - \tilde{C}$, then

$$\begin{aligned} \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} A \widehat{\beta}^{(-k)} - \beta^{*\top} A \beta^*)^2 &= \mathbb{E}_{I_k^c}(\|\widehat{\beta}^{(-k)} - \beta^*\|_A^2 + 2(\widehat{\beta}^{(-k)} - \beta^*)^\top A \beta^*)^2 \\ &\leq 2\mathbb{E}_{I_k^c}\|\widehat{\beta}^{(-k)} - \beta^*\|_A^4 + 8\mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top A \beta^*)^2 \\ &= \mathcal{O}(\mathbb{E}_{I_k^c}\|\widehat{\beta}^{(-k)} - \beta^*\|_A^4 + (\mathbb{E}_{I_k^c}\|\widehat{\beta}^{(-k)} - \beta^*\|_A^4 \mathbb{E}\|\beta^*\|_A^4)^{\frac{1}{2}}) \\ &= o(1), \end{aligned}$$

$$\begin{aligned} \mathbb{E}_{I_k^c}\|\widehat{\beta}^{(-k)} - \beta^*\|_A^4 &= \mathcal{O}(\mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)} - \beta^*)^\top \tilde{V})^4 + \|\widehat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}^4) \\ &= \mathcal{O}(\|\widehat{\beta}^{(-k)} - \beta^*\|_{\tilde{C}}^4) = o_P(1), \end{aligned}$$

$$\begin{aligned} \mathbb{E}\|\beta^*\|_A^4 &= \text{Var}(\beta^{*\top} \tilde{V})^2 \leq \mathbb{E}(\beta^{*\top} \tilde{V})^4 \\ &\leq b^4 \sup_{\|a\|_2=1} \mathbb{E}(a^\top Z)^4 = \mathcal{O}(1). \end{aligned}$$

By the fact that $\mathbb{E}a^2 = \mathbb{E}b^2 + \mathbb{E}(a-b)(a+b) = \mathbb{E}b^2 + \mathcal{O}((\mathbb{E}(a-b)^2 \mathbb{E}(a^2 + b^2))^{\frac{1}{2}})$,

$$\begin{aligned} \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\widehat{V}\widehat{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)})^2 &= \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\tilde{V}\tilde{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)})^2 \\ &\quad + \mathcal{O}_P((\mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)\top} \widehat{V})^2 - (\widehat{\beta}^{(-k)\top} \tilde{V})^2))^{\frac{1}{2}}), \end{aligned}$$

$$\begin{aligned} \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\tilde{V}\tilde{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)})^2 &= \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \widehat{\beta}^{(-k)})^2 \\ &\quad + \mathcal{O}_P(\widehat{\beta}^{(-k)\top} (\widehat{C} - \tilde{C}) \widehat{\beta}^{(-k)}) \\ &= \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \widehat{\beta}^{(-k)})^2 + o_P(1), \end{aligned}$$

$$\begin{aligned} \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \widehat{\beta}^{(-k)})^2 &= \mathbb{E}_{I_k^c}(\beta^{*\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \beta^*)^2 \\ &\quad + \mathcal{O}_P((\mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} A \widehat{\beta}^{(-k)} - \beta^{*\top} A \beta^*)^2)^{\frac{1}{2}}) \\ &= \mathbb{E}_{I_k^c}(\beta^{*\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \beta^*)^2 + o_P(1). \end{aligned}$$

By the fact that $\mathbb{E}((a^2 - b^2))^2 = \mathbb{E}(a-b)^2(a+b)^2 = \mathcal{O}((\mathbb{E}(a-b)^4 \mathbb{E}(a+b)^4)^{\frac{1}{2}})$,

$$\mathbb{E}_{I_k^c}((\widehat{\beta}^{(-k)\top} \widehat{V})^2 - (\widehat{\beta}^{(-k)\top} \tilde{V})^2)^2 = \mathcal{O}_P((\widehat{\beta}^{(-k)\top} (\widehat{\mu} - \tilde{\mu}))^2) = o_P(1).$$

Combining all the previous results,

$$\begin{aligned} N^{-1} \sum_{i \in I_k} \xi_i^{(k)2} &= \mathbb{E}_{I_k^c}(\widehat{\beta}^{(-k)\top} (\widehat{V}\widehat{V}^\top - \widehat{C}) \widehat{\beta}^{(-k)})^2 + o_P(1) \\ &= \mathbb{E}(\beta^{*\top} (\tilde{V}\tilde{V}^\top - \tilde{C}) \beta^*)^2 + o_P(1). \end{aligned}$$

Similarly,

$$(60) \quad N^{-1} \sum_{i \in I_k} \nu_i^{(k)2} = N^{-1} \sum_{i \in I_k} (\nu_i^{(k)} - \mathbb{E}_{I_k^c} \nu_i^{(k)})^2 + (N^{-1} \sum_{i \in I_k} \nu_i^{(k)} - \mathbb{E}_{I_k^c} \nu_i^{(k)})^2$$

$$(61) \quad = \text{Var}_{P_X}(\widehat{\varepsilon}^2 + 2 \widehat{\beta}^{(-k)\top} \widehat{V} \widehat{\varepsilon})^2 + o_P(1)$$

$$(62) \quad = \text{Var}_{P_X}(\varepsilon^2 + 2 \beta^{*\top} \widetilde{V} \varepsilon)^2 + o_P(1).$$

Therefore,

$$\widehat{\sigma}_\nu^2 + \frac{n}{m} \widehat{\sigma}_\xi^2 = \text{Var}(\varepsilon^2 + 2 \beta^{*\top} \widetilde{V} \varepsilon) + \tau \text{Var}(\beta^{*\top} \widetilde{V})^2 + o_P(1).$$

Proof of Theorem 6. This proof provides consistency rates of $\widehat{\delta}$ and a asymptotic normal result for $\sqrt{n}(\widehat{\delta} - \delta)$ when some specific rates are satisfied. Assume $a_{n,p} = \mathcal{O}(1)$, $b_{m,p} \mathcal{O}(1)$, $a_{n,p} b_{m,p} = \mathcal{O}(n^{-\frac{1}{2}})$. Let

$$r_i^{(-k)} = D_i / \widehat{e}^{(-k)}(X_i), \quad r_i = D_i / e(X_i), \quad r = D / e(X)$$

then

$$\mathbb{E}_{I_k^c} (r^{(-k)} - r)^2 = \mathbb{E}_{I_k^c} \frac{D(\widehat{e}^{(-k)}(X) - e(X))^2}{(\widehat{e}^{(-k)}(X)e(X))^2} = \mathcal{O}_P(b_{m,p}^2).$$

By the definition of $\widehat{\tau}_1^{(k)}$ as in (29),

$$\begin{aligned} \widehat{\tau}_1^{(k)} &= \beta_1^{*\top} \widehat{\mu} + N^{-1} \sum_{i \in I_k} r_i (Y_i - \beta_1^{*\top} \widetilde{X}_i) + (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widehat{\mu} + N^{-1} \sum_{i \in I_k} (r_i^{(-k)} - r_i) (Y_i - \beta_1^{*\top} \widetilde{X}_i) \\ &\quad - N^{-1} \sum_{i \in I_k} r_i (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X}_i - N^{-1} \sum_{i \in I_k} (r_i^{(-k)} - r_i) (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X}_i. \end{aligned}$$

By Conditions 1, 5, 6 and Lemma 2,

$$\begin{aligned} \beta_1^{*\top} \widehat{\mu} &= \beta_1^{*\top} \widetilde{\mu} + \mathcal{O}_P(m^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} r_i (Y_i - \beta_1^{*\top} \widetilde{X}_i) &= \mathbb{E} \varepsilon \zeta / e(X) + \mathcal{O}_P(N^{-\frac{1}{2}}) = \mathcal{O}_P(N^{-\frac{1}{2}}), \\ (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widehat{\mu} &= (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{\mu} + \mathcal{O}_P(a_{n,p} m^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} (r_i^{(-k)} - r_i) (Y_i - \beta_1^{*\top} \widetilde{X}_i) &= \mathbb{E}_{I_k^c} (r^{(-k)} - r) \varepsilon + \mathcal{O}_P(b_{m,p} N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} r_i (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X}_i &= (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{\mu} + \mathbb{E}_{I_k^c} \frac{\zeta (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X}}{e(X)} + \mathcal{O}_P(a_{n,p} N^{-\frac{1}{2}}) \\ &= (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{\mu} + \mathcal{O}_P(a_{n,p} N^{-\frac{1}{2}}), \\ N^{-1} \sum_{i \in I_k} (r_i^{(-k)} - r_i) (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X}_i &= \mathbb{E}_{I_k^c} (r^{(-k)} - r) (\widehat{\beta}_1^{(-k)} - \beta_1^*)^\top \widetilde{X} + \mathcal{O}_P(a_{n,p} b_{m,p} N^{-\frac{1}{2}}) \\ &= \mathcal{O}_P(a_{n,p} b_{m,p}). \end{aligned}$$

Hence,

$$\widehat{\tau}_1^{(k)} = \beta_1^{*\top} \widetilde{\mu} + \mathcal{O}_P(a_{n,p} b_{m,p} + (1 + a_{n,p} + b_{m,p}) N^{-\frac{1}{2}}).$$

Similarly,

$$\widehat{\tau}_2^{(k)} = \beta_2^{*\top} \widetilde{\mu} + \mathcal{O}_P(a_{n,p} b_{m,p} + (1 + a_{n,p} + b_{m,p}) N^{-\frac{1}{2}}).$$

When $K < \infty$,

$$\widehat{\delta} = \widehat{\tau}_1 - \widehat{\tau}_2 = \delta + \mathcal{O}_P(a_{n,p} b_{m,p} + (1 + a_{n,p} + b_{m,p}) n^{-\frac{1}{2}}) = \delta + \mathcal{O}_P(n^{-\frac{1}{2}}).$$

Moreover, if $a_{n,p} = o_P(1)$, $b_{m,p} = o_P(1)$, $a_{n,p} b_{m,p} = o_P(n^{-\frac{1}{2}})$. Then, by the previous results and Lindeberg-Feller Central Limit Theorem,

$$\begin{aligned} \sqrt{n}(\widehat{\delta} - \delta) &= \sqrt{n}(\beta_1^* - \beta_0^*)^\top (\widehat{\mu} - \mu) + n^{-\frac{1}{2}} \sum_{i=1}^n \varepsilon_i \zeta_i / (e(X_i)(1 - e(X_i))) + o_P(1) \\ &\xrightarrow{d} \mathcal{N} \left(0, \text{Var} \left(\frac{\varepsilon \zeta}{e(X)(1 - e(X))} \right) + \tau(\beta_1^* - \beta_0^*)^\top \widetilde{C}(\beta_1^* - \beta_0^*) \right). \end{aligned}$$

Proof of Theorem 7. This proof provides a consistency result for \widehat{V}_δ . Similarly as in the Proof of Theorem 5,

$$\begin{aligned} \widehat{\sigma}_{\nu,\delta}^{2(k)} &= N^{-1} \sum_{i \in I_k} (\nu_{\delta,i}^{(k)} - \mathbb{E}_{I_k^c} \nu_{\delta}^{(k)})^2 + (N^{-1} \sum_{i \in I_k} \nu_{\delta,i}^{(k)} - \mathbb{E}_{I_k^c} \nu_{\delta}^{(k)})^2 \\ &= \text{Var}_{P_{X,\varepsilon,\zeta}}(r^{(-k)}(Y - \widehat{\beta}_1^{(-k)\top} \widetilde{X}) - \rho^{(-k)}(Y - \widehat{\beta}_0^{(-k)\top} \widetilde{X})) + o_P(1) \\ &= \text{Var}_{P_{X,\varepsilon,\zeta}}(r(Y - \beta_1^{*\top} \widetilde{X}) - \rho(Y - \beta_0^{*\top} \widetilde{X})) + o_P(1) \\ &= \text{Var} \left(\frac{\varepsilon \zeta}{e(X)(1 - e(X))} \right) + o_P(1) \end{aligned}$$

and

$$\begin{aligned} \widehat{\sigma}_{\xi,\delta}^{2(k)} &= \mathbb{E}_{I_k^c}((\widehat{\beta}_1^{(-k)} - \widehat{\beta}_0^{(-k)})^\top (\widetilde{X} - \widehat{\mu}))^2 + o_P(1) \\ &= \mathbb{E}_{I_k^c}((\beta_1^* - \beta_0^*)^\top (\widetilde{X} - \widetilde{\mu}))^2 + o_P(1) \\ &= (\beta_1^* - \beta_0^*)^\top \widetilde{C}(\beta_1^* - \beta_0^*) + o_P(1). \end{aligned}$$

Therefore,

$$\widehat{V}_\delta = V_\delta + o_P(1).$$

Proof of Theorem 8. This proof provides an asymptotic normal result for $\sqrt{n}(\widehat{d} - d)$. Let $\mathcal{A} = \{\widehat{\sigma}_Y \neq 0\}$, then for any $c \in (0, \sigma^2)$,

$$\mathbb{P}(\mathcal{A}) \geq \mathbb{P}(|\widehat{\sigma}^2 - \sigma^2| < c) = 1 + o(1)$$

since $\widehat{\sigma}^2 = \sigma^2 + o_P(1)$. Now conditional on event \mathcal{A} . Define

$$\begin{aligned} \nu_i^* &= \varepsilon_i^2 - \sigma_\varepsilon^2 + 2\beta^{*\top} \widetilde{V}_i \varepsilon, & \xi_i^* &= (\beta^{*\top} \widetilde{V}_i)^2 - b^2, \\ \nu_{\delta,i}^* &= \frac{D_i \varepsilon_i}{e(X_i)} - \frac{(1 - D_i) \varepsilon_i}{1 - e(X_i)}, & \xi_{\delta,i}^* &= (\beta_1^* - \beta_0^*)^\top \widetilde{V}_i. \end{aligned}$$

From Theorems 5 and 6,

$$\begin{aligned}\sqrt{n}(\widehat{\sigma}^2 - \sigma^2) &= n^{-\frac{1}{2}} \sum_{i=1}^n \nu_i^* + n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} \xi_i^* + o_P(1), \\ \sqrt{n}(\widehat{\delta} - \delta) &= n^{-\frac{1}{2}} \sum_{i=1}^n \nu_{\delta,i}^* + n^{\frac{1}{2}} m^{-1} \sum_{i=n+1}^{m+n} \xi_{\delta,i}^* + o_P(1).\end{aligned}$$

Since

$$\mathbb{E}|\nu_i^*|^{2+c} < c, \quad \mathbb{E}|\xi_i^*|^{2+c} < c, \quad \mathbb{E}|\nu_{\delta,i}^*|^{2+c} < c, \quad \mathbb{E}|\xi_{\delta,i}^*|^{2+c} < c,$$

by Lindeberg-Feller Central Limit Theorem,

$$\begin{aligned}\sqrt{n}(\widehat{d} - d) &= \sqrt{n} \left(\frac{\widehat{\delta}}{\widehat{\sigma}} - \frac{\delta}{\sigma} \right) = \sqrt{n} \left(\frac{\widehat{\delta}\sigma - \delta\widehat{\sigma}}{\sigma\widehat{\sigma}} \right) \\ &= \frac{\sqrt{n}(\widehat{\delta} - \delta)}{\widehat{\sigma}} - \frac{\delta\sqrt{n}(\widehat{\sigma}^2 - \sigma^2)}{\sigma\widehat{\sigma}(\widehat{\sigma} + \sigma)} \\ &= \frac{n^{-\frac{1}{2}} \sum_{i=1}^n \nu_{\delta,i}^* + \tau m^{-\frac{1}{2}} \sum_{i=n+1}^{m+n} \xi_{\delta,i}^*}{\sigma} - \frac{\delta(n^{-\frac{1}{2}} \sum_{i=1}^n \nu_i^* + \tau m^{-\frac{1}{2}} \sum_{i=n+1}^{m+n} \xi_i^*)}{2\sigma^3} + o_P(1) \\ &= n^{-\frac{1}{2}} \sum_{i=1}^n \left(\frac{\nu_{\delta,i}^*}{\sigma} - \frac{\delta\nu_i^*}{2\sigma^3} \right) + \tau m^{-\frac{1}{2}} \sum_{i=n+1}^{m+n} \left(\frac{\xi_{\delta,i}^*}{\sigma} - \frac{\delta\xi_i^*}{2\sigma^3} \right) + o_P(1) \\ &\xrightarrow{d} \mathcal{N}(0, V_d),\end{aligned}$$

where

$$V_d = \text{Var} \left(\frac{\nu_{\delta,i}^*}{\sigma} - \frac{\delta\nu_i^*}{2\sigma^3} \right) + \tau \text{Var} \left(\frac{\xi_{\delta,i}^*}{\sigma} - \frac{\delta\xi_i^*}{2\sigma^3} \right).$$

Proof of Theorem 9. This proof provides a consistency result for \widehat{V}_d . By the definition of $\nu_{d,i}$,

$$\begin{aligned}N^{-1} \sum_{i \in I_k} \nu_{d,i}^2 &= N^{-1} \sum_{i \in I_k} \left(\frac{1}{\widehat{\sigma}} \nu_{\delta,i} - \frac{\widehat{\delta}}{2\widehat{\sigma}^3} \nu_i \right)^2 \\ &= N^{-1} \sum_{i \in I_k} \left(\frac{1}{\sigma} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \nu_i \right)^2 + 2 \left(\frac{1}{\widehat{\sigma}} - \frac{1}{\sigma} \right) N^{-1} \sum_{i \in I_k} \nu_{\delta,i} \left(\frac{1}{\sigma} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \nu_i \right) \\ &\quad - 2 \left(\frac{\widehat{\delta}}{2\widehat{\sigma}^3} - \frac{\delta}{2\sigma^3} \right) N^{-1} \sum_{i \in I_k} \nu_i \left(\frac{1}{\sigma} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \nu_i \right) \\ &= N^{-1} \sum_{i \in I_k} \left(\frac{1}{\sigma} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \nu_i \right)^2 + o_P(1).\end{aligned}$$

By the similar approach as (62),

$$\begin{aligned}
& N^{-1} \sum_{i \in I_k} \left(\frac{1}{\sigma} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \nu_i \right)^2 \\
&= N^{-1} \sum_{i \in I_k} \left(\frac{1}{\sigma} (\nu_{\delta,i} - \mathbb{E}_{I_k^c} \nu_{\delta,i}) - \frac{\delta}{2\sigma^3} (\nu_i - \mathbb{E}_{I_k^c} \nu_i) \right)^2 - \left(\frac{1}{\sigma} \mathbb{E}_{I_k^c} \nu_{\delta,i} - \frac{\delta}{2\sigma^3} \mathbb{E}_{I_k^c} \nu_i \right)^2 \\
&= \text{Var}_{I_k^c} \left(\frac{\nu_{\delta,i}}{\sigma} - \frac{\delta \nu_i}{2\sigma^3} \right) + o_P(1) \\
&= \text{Var}_{I_k^c} \left(\frac{r_i^{(-k)} (Y_i - \widehat{\beta}_1^{(-k)\top} \widetilde{X}_i) - \rho_i^{(-k)} (Y_i - \widehat{\beta}_0^{(-k)\top} \widetilde{X}_i)}{\sigma} \right. \\
&\quad \left. - \frac{\delta ((Y_i - \delta)^2 - (\widehat{\beta}_{\widetilde{W}}^{(-k)\top} (\widetilde{W}_i - \widetilde{\mu}_{\widetilde{W}}))^2)}{2\sigma^3} \right) + o_P(1) \\
&= \text{Var} \left(\frac{r_i (Y_i - \beta_1^{*\top} \widetilde{X}_i) - \rho_i (Y_i - \beta_0^{*\top} \widetilde{X}_i)}{\sigma} \right. \\
&\quad \left. - \frac{\delta ((Y_i - \delta)^2 - (\beta_{\widetilde{W}}^{*\top} (\widetilde{W}_i - \widetilde{\mu}_{\widetilde{W}}))^2)}{2\sigma^3} \right) + o_P(1) \\
&= \text{Var} \left(\frac{\nu_{\delta,i}^*}{\sigma} - \frac{\delta \nu_i^*}{2\sigma^3} \right) + o_P(1).
\end{aligned}$$

Similarly,

$$N^{-1} \sum_{i \in I_k} \left(\frac{\xi_{\delta,i}}{\widehat{\sigma}} - \frac{\widehat{\delta} \xi_i}{2\widehat{\sigma}^3} \right)^2 = \text{Var} \left(\frac{\xi_{\delta,i}^*}{\sigma} - \frac{\delta \xi_i^*}{2\sigma^3} \right) + o_P(1).$$

Therefore,

$$\widehat{V}_d = V_d + o_P(1).$$

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA SAN DIEGO
E-mail address: yuz643@ucsd.edu; jbradic@ucsd.edu