

Semi-Supervised Quantile Estimation: Robust and Efficient Inference in High Dimensional Settings

Abhishek Chakrabortty*

Guorong Dai†

Raymond J. Carroll‡

August 15, 2024

Abstract

We consider quantile estimation in a semi-supervised setting, where one has two available data sets: (i) a small or moderate sized *labeled data set* containing observations for a response and a set of possibly high dimensional covariates, and (ii) a much larger *unlabeled data set* where only the covariates are observed. Such settings are of increasing relevance in modern studies involving large databases where labeled data may be limited due to practical constraints but unlabeled data are plentiful, and it is of interest to investigate how the latter may be exploited. We propose a family of semi-supervised estimators for the response quantile(s) based on the two data sets, to improve the estimation accuracy compared to the supervised estimator, i.e., the sample quantile, which uses the labeled data only. These estimators are based on a flexible imputation strategy applied to the estimating equation along with a debiasing step that allows for full robustness against misspecification of the imputation model. Further, a one-step update strategy is adopted to enable easy implementation of our method and handle the inevitable complexity arising from the non-linear nature of the quantile estimating equation. Under fairly mild assumptions, we prove our estimators are *fully robust* to the choice of the nuisance imputation model, in the sense of *always* maintaining root- n consistency and asymptotic normality, while having *improved* efficiency relative to the supervised estimator. Further, they achieve semi-parametric *optimality* also, provided the relation between the response and the covariates is correctly specified via the imputation model. In addition, as an illustration of estimating the nuisance imputation function, we consider kernel smoothing type estimators on lower dimensional and possibly estimated transformations of the high dimensional covariates, and we establish novel results on *uniform* convergence rates of such kernel smoothing estimators in *high dimensions*, involving responses indexed by a function class and usage of dimension reduction techniques. These results may be of independent interest. Numerical results on both simulated and real data confirm our semi-supervised approach's improved performance, both in terms of estimation as well as inference.

Keywords: Semi-supervised inference, Quantile estimation, Robustness and efficiency, Imputation and debiasing, High dimensional nuisance estimators, Kernel smoothing with dimension reduction.

*Department of Statistics, Texas A&M University (email: abhishek@stat.tamu.edu).

†Department of Statistics and Data Science, Fudan University (email: guorongdai@fudan.edu.cn). Guorong Dai is a joint first author. Guorong Dai is the corresponding author.

‡Department of Statistics, Texas A&M University (email: carroll@stat.tamu.edu).

1 Introduction

Semi-supervised settings, as the name suggests, are characterized by two available data sets: (i) a *labeled data* containing a limited number of observations on both a response Y and a set of covariates \mathbf{X} , and (ii) an *unlabeled data* of much larger size where only the covariates \mathbf{X} are observed. Statistical learning in such settings, often termed “semi-supervised learning”, has gained substantial attention in the last two decades in machine learning, and more recently, in the statistics community as well. A detailed overview of semi-supervised learning and the growing literature can be found in [Zhu and Goldberg \(2009\)](#) and [Chapelle et al. \(2010\)](#). Semi-supervised settings emerge naturally when observations for \mathbf{X} are easy to collect for a large cohort, but obtaining the corresponding Y is costly or time-consuming. This situation is ubiquitous in modern studies across scientific disciplines (see [Chapelle et al. \(2010\)](#) for various examples), including modern biomedical studies, such as electronic health records – rich resources of data for discovery research – where labeling of Y is often logically prohibitive; see Section 1 of [Chakrabortty and Cai \(2018\)](#) for further details. Another highly relevant biomedical application is in integrative genomics, especially expression quantitative trait loci studies ([Michaelson et al., 2009](#)) that develop association mapping between gene expression levels Y and genetic variants \mathbf{X} . A major bottleneck for the power of such studies has turned out to be the limited sample size of the expensive gene expression data ([Flutre et al., 2013](#); [McCaw et al., 2021](#)). In contrast, observations for the genetic variants are usually cheaper and available for much more individuals. Safely and efficiently leveraging these plentiful unlabeled data thus necessitates developing suitable semi-supervised strategies for robust and efficient inference, which are our main focus in this paper.

1.1 Our goal and existing work

In semi-supervised settings, the most critical and natural question is when and how the unlabeled data can be used to improve estimation accuracy compared to *supervised* methods, which take account of the labeled data only. In principle, such improvement is determined by how the parameter of interest depends on the marginal distribution $\mathbb{P}_{\mathbf{X}}$ of \mathbf{X} as the unlabeled data are informative for $\mathbb{P}_{\mathbf{X}}$ only ([Seeger, 2002](#); [Zhang and Oles, 2000](#)). Thus conditions on the relation between $\mathbb{P}_{\mathbf{X}}$ and the conditional distribution of Y given \mathbf{X} were assumed explicitly or implicitly by many existing semi-supervised approaches that aimed to estimate $\mathbb{E}(Y | \mathbf{X})$ nonparametrically, including generative modeling ([Nigam et al., 2000](#)), graph-based methods ([Zhu, 2005](#)) and manifold regularization ([Belkin et al., 2006](#)). Their advantages over supervised learning, however, cannot be guaranteed when these distributional assumptions are violated ([Cozman and Cohen, 2001](#); [Cozman et al., 2003](#)). Recently, robust semi-supervised methods were proposed for various (finite dimensional) *inference* problems, including mean estimation ([Zhang et al., 2019](#); [Zhang and Bradic, 2022](#)), linear regression ([Chakrabortty and Cai, 2018](#); [Azriel et al., 2022](#)) and prediction accuracy evaluation ([Gronsbell and Cai, 2018](#)), among others. As the name “robust” suggests, they have been proven to be at least as efficient as their supervised counterparts even if the underlying distribution or model assumptions are misspecified. In a more general context, Chapter 2 of [Chakrabortty \(2016\)](#) provides an elaborate interpretation about robustness and efficiency of semi-supervised inference from a semi-parametric perspective. Despite this rich and growing literature on semi-supervised inference, the problem of *quantile estimation* has, to the best of our knowledge, not been investigated in semi-supervised settings. Quantile(s), however, are important parameter(s) that help characterize the whole response distribution, and are often robust choices for measuring the central

tendency of heavy-tailed and/or skewed data, compared to the mean (which may not even exist). This work therefore aims to bridge this gap in the existing semi-supervised learning literature regarding inference for quantiles.

1.2 Problem setup

To formulate our work, suppose $\mathbf{X} \in \mathcal{X} \subset \mathbb{R}^p$, for some \mathcal{X} , and $Y \in \mathbb{R}$ is continuous with a density function $f(\cdot)$ as well as a distribution function $F(\cdot)$. Here, the dimension $p \equiv p_n$ of \mathbf{X} could *diverge*, and possibly exceed the sample size n , allowing for the case where $p \gg n$. Our goal is to estimate the τ -quantile, $\theta_0 \equiv \theta_0(\tau) \in \Theta \subset \mathbb{R}$, for some Θ , of the response Y , for some fixed and known $\tau \in (0, 1)$, defined as the solution to the estimating equation:

$$\mathbb{E}\{\psi(Y, \theta_0)\} = 0, \quad (1.1)$$

where $\psi(y, \theta) := I(y < \theta) - \tau$ with $I(\cdot)$ being the indicator function. It is important to notice that the existing robust semi-supervised methods for the problems as mentioned in the last paragraph *cannot* be easily extended to this problem, because they all directly impute the unlabeled Y with the conditional mean $\mathbb{E}(Y | \mathbf{X})$, or relevant working models thereof. Due to the non-linear and inseparable nature of Y and θ_0 in the equation (1.1), such imputation strategies can, however, lead to *bias* in this case. This fact makes quantile estimation fundamentally different and much more challenging compared to problems based on estimating equations that are linear/separable in the response Y and the parameter(s) of interest, such as mean estimation (Zhang et al., 2019; Zhang and Bradic, 2022) and linear regression (Chakrabortty and Cai, 2018; Azriel et al., 2022).

In a semi-supervised setting, we have two independent data sets: the labeled data $\mathcal{L} := \{\mathbf{Z}_i := (Y_i, \mathbf{X}_i^T)^T : i = 1, \dots, n\}$ and the unlabeled data $\mathcal{U} := \{\mathbf{X}_i : i = n + 1, \dots, n + N\}$, which consist of n and N independent copies of $\mathbf{Z} := (Y, \mathbf{X}^T)^T$ and \mathbf{X} , respectively. Here, the labeling mechanism is typically by design, so that \mathcal{L} and \mathcal{U} have the same distribution – a standard, often implicit, assumption in the literature; see Chakrabortty and Cai (2018, Remark 2.1) for details – and the labeling status, i.e., labeled or not, of an individual is considered fixed/non-random, unlike usual missing data problems where the missingness mechanism/indicator is considered random. More importantly, a major challenge and a key feature – often a consequence of the underlying practical circumstances – contrasting our setup from traditional missing data problems is that the unlabeled data can be *much larger* in size than the labeled data ($N \gg n$), i.e., it is possible that

$$\lim_{n, N \rightarrow \infty} (n/N) = 0.$$

An example is the *ideal semi-supervised setting* where n is finite and $N = \infty$, i.e., the distribution $\mathbb{P}_{\mathbf{X}}$ is known. This clearly violates the “positivity assumption”, that the proportion of Y observed in the whole data $\mathcal{L} \cup \mathcal{U}$ is bounded away from zero, often required in the missing data literature (Tsiatis, 2007; Little and Rubin, 2019). Thus, letting $\nu_{n, N} := n/(n + N)$, we allow

$$\nu := \lim_{n, N \rightarrow \infty} \nu_{n, N} \in [0, 1].$$

This natural violation of the positivity assumption caused by the case $\nu = 0$ in fact raises substantial technical difficulties, e.g., non-standard asymptotics, that cannot be handled by the classical theory of missing data, making our semi-supervised setting fairly unique and challenging.

1.3 Our contributions

Based on the labeled data \mathcal{L} only, a *supervised* estimator $\hat{\theta}_{\text{SUP}}$ can be obtained by solving the sample version of (1.1) given by $n^{-1}\sum_{i=1}^n\{I(Y_i < \hat{\theta}_{\text{SUP}}) - \tau\} = 0$. Although $\hat{\theta}_{\text{SUP}}$ possesses consistency and asymptotic normality, which will be shown in Proposition 2.1, ignoring the unlabeled data \mathcal{U} generally leads to a loss of efficiency; see Remark 2.5 for details. To improve the estimation of θ_0 , this article proposes a family of semi-supervised estimators, which takes both \mathcal{L} and \mathcal{U} into consideration by imputing the function $\psi(Y, \theta)$ in the definition (1.1). Under mild conditions, our estimators are *always $n^{1/2}$ -consistent for θ_0 and asymptotically normal* with arbitrary imputation functions (Theorem 2.1), and also ensured to *outperform* the supervised estimator $\hat{\theta}_{\text{SUP}}$ with respect to asymptotic variance (Remark 2.5). When the imputation function correctly specifies the conditional distribution of Y given \mathbf{X} , our method further attains *semi-parametric efficiency* (Remark 2.5). Another advantage of our estimators is their *first-order insensitivity*, that is, their influence functions are not affected by estimation errors or knowledge of the nuisance estimator's construction (Remark 2.4). This property is particularly desirable for inference when the covariates \mathbf{X} are high dimensional or the nuisance functions are estimated by nonparametric techniques. Our method uses sample splitting/cross-fitting techniques to allow for such nuisance estimators. In the proof of these claimed properties, technical innovations related to empirical process theory are established to handle the classes of $\psi(Y, \theta)$ and other relevant random functions indexed by θ ; see Lemma A.2 in Appendix A.2. When constructing our estimators, we adopt the strategy of one-step update (Van der Vaart, 2000; Tsiatis, 2007) to overcome computational difficulties arising from the inseparability of Y and θ in (1.1), providing a simple implementation of our estimation and inference procedures. It also avoids the burdensome task of estimating nuisance functions for the whole parameter space Θ . Instead, we only need to consider one single value of θ . This feature allows us to use a wide range of approaches, including parametric regression and nonparametric/machine learning approaches, like kernel smoothing and random forest, for the nuisance estimation, a fairly important component in the implementation of our method that poses substantial challenges in high dimensional scenarios. Although the above-mentioned desirable properties of our method are guaranteed by *any* nuisance estimators as long as the high-level conditions in Theorem 2.1 hold, we thoroughly study *kernel smoothing estimators*, with possible use of *dimension reduction*, as an illustration. Specifically, we show those high-level conditions are satisfied by kernel smoothing estimators through deriving their uniform convergence rates, when the outcome involves a function class of Y and the covariates are generated by transforming the high dimensional \mathbf{X} via possibly unknown dimension reduction processes; see Theorem 3.1 and Remark 3.1. These results extend the existing theory of kernel smoothing estimators with generated covariates (Mammen et al., 2012; Escanciano et al., 2014; Mammen et al., 2016) in high dimensions. These are also useful in other applications, and should be of independent interest. In summary, our main contributions thus are:

- (a) We develop a globally robust and locally optimal strategy for quantile estimation in semi-supervised and high dimensional setups, providing a family of $n^{1/2}$ -consistent, asymptotically normal and first-order insensitive estimators ensured to be *at least as efficient* as the sample quantile $\hat{\theta}_{\text{SUP}}$ and more efficient whenever possible; see Theorem 2.1 and Remarks 2.4–2.5;
- (b) As an illustration of the nuisance estimators and their theoretical properties required in our method, we consider kernel smoothing type estimators, and derive their uniform convergence rates when the response is indexed by a function class and the high dimensional covariates are transformed by (unknown) dimension reduction mechanisms; see Theorem 3.1 and Remark 3.1.

1.4 Organization of the article

In the next section, we introduce our family of semi-supervised estimators for the response quantile θ_0 and study their asymptotic properties. The choice and estimation of the nuisance functions involved in our approach are theoretically investigated in Section 3. Section 4 provides numerical results from extensive simulations under a wide range of data generating mechanisms, followed by an empirical data example in Section 5. These numerical results substantiate the properties and advantages of our approach stated in the previous sections. Finally, Section 6 ends the article with a concluding remark as well as a brief discussion of possible future work. All technical details, including auxiliary lemmas and proofs of all theoretical results, and extra numerical results as well as necessary supplements to the data analysis in Section 5 can be found in Appendices A–C.

2 Semi-supervised estimation of quantiles via one-step update

Notation. Throughout, we use the lower case letter c to represent a generic positive constant, including c_1, c_2 , etc., which may vary from line to line. For a $d_1 \times d_2$ matrix \mathbf{P} whose (i, j) th component is $\mathbf{P}_{[ij]}$, we define $\|\mathbf{P}\|_0 := \max_{1 \leq j \leq d_2} \{\sum_{i=1}^{d_1} I(\mathbf{P}_{[ij]} \neq 0)\}$, $\|\mathbf{P}\|_1 := \max_{1 \leq j \leq d_2} (\sum_{i=1}^{d_1} |\mathbf{P}_{[ij]}|)$, $\|\mathbf{P}\| := \max_{1 \leq j \leq d_2} \{(\sum_{i=1}^{d_1} \mathbf{P}_{[ij]}^2)^{1/2}\}$ and $\|\mathbf{P}\|_\infty := \max_{1 \leq i \leq d_1, 1 \leq j \leq d_2} |\mathbf{P}_{[ij]}|$. The symbols $\mathbf{1}_d$ and $\mathbf{0}_d$ refer to d -dimensional vectors of ones and zeros, respectively and $Normal(\mu, \sigma^2)$ denotes the Normal distribution with mean μ and variance σ^2 . We denote $\mathcal{B}(\alpha, \varepsilon) := \{a : |a - \alpha| \leq \varepsilon\}$ as a generic neighborhood of a scalar α with some radius $\varepsilon > 0$. For a vector $\boldsymbol{\beta}$, we use $\boldsymbol{\beta}_{[j]}$ to denote its j th component. For any random function $\hat{g}(\cdot, \theta)$ and a random vector \mathbf{W} with copies $\mathbf{W}_1, \dots, \mathbf{W}_{n+N}$, we denote $\mathbb{E}_{\mathbf{W}}\{\hat{g}(\mathbf{W}, \theta)\} := \int \hat{g}(\mathbf{w}, \theta) dF_{\mathbf{W}}(\mathbf{w})$ as the expectation of $\hat{g}(\mathbf{W}, \theta)$ with respect to \mathbf{W} treating $\hat{g}(\cdot, \theta)$ as a non-random function, where $F_{\mathbf{W}}(\cdot)$ is the distribution function of \mathbf{W} . For $M \in \{n, n+N\}$, we write $\mathbb{E}_M\{\hat{g}(\mathbf{W}, \theta)\} := M^{-1} \sum_{i=1}^M \hat{g}(\mathbf{W}_i, \theta)$ and $\mathbb{G}_M\{\hat{g}(\mathbf{W}, \theta)\} := M^{1/2} [\mathbb{E}_M\{\hat{g}(\mathbf{W}, \theta)\} - \mathbb{E}_{\mathbf{W}}\{\hat{g}(\mathbf{W}, \theta)\}]$ as well as define $\text{var}_M\{\hat{g}(\mathbf{W}, \theta)\} := \mathbb{E}_M[\{\hat{g}(\mathbf{W}, \theta)\}^2] - [\mathbb{E}_M\{\hat{g}(\mathbf{W}, \theta)\}]^2$. Also, we let $\mathbb{E}_N\{\hat{g}(\mathbf{W}, \theta)\} := N^{-1} \sum_{i=n+1}^{n+N} \hat{g}(\mathbf{W}_i, \theta)$ and $\mathbb{G}_N\{\hat{g}(\mathbf{W}, \theta)\} := N^{1/2} [\mathbb{E}_N\{\hat{g}(\mathbf{W}, \theta)\} - \mathbb{E}_{\mathbf{W}}\{\hat{g}(\mathbf{W}, \theta)\}]$. Lastly, let $f(\cdot | \mathbf{w})$ and $F(\cdot | \mathbf{w})$ represent the conditional density and distribution functions of Y given $\mathbf{W} = \mathbf{w}$, respectively.

2.1 The supervised estimator

We first investigate the supervised estimator $\hat{\theta}_{\text{SUP}}$, i.e., the sample quantile, solving: $\mathbb{E}_n\{I(Y < \hat{\theta}_{\text{SUP}}) - \tau\} = 0$. To study its limiting behavior, we need the following basic assumption.

Assumption 2.1. The parameter θ_0 is in the interior of the space Θ . The density function $f(\cdot)$ of Y satisfies $f(\theta_0) > 0$ and has a bounded derivative in $\mathcal{B}(\theta_0, \varepsilon)$.

The basic Assumption 2.1 is fairly standard, and it guarantees the identifiability of θ_0 as well as the validity of $\hat{\theta}_{\text{SUP}}$. The following proposition then gives the limiting properties of $\hat{\theta}_{\text{SUP}}$.

Proposition 2.1. *Under Assumption 2.1, the supervised estimator $\hat{\theta}_{\text{SUP}}$ satisfies:*

$$\hat{\theta}_{\text{SUP}} - \theta_0 = -\{nf(\theta_0)\}^{-1} \sum_{i=1}^n \psi(Y, \theta_0) + o_p(n^{-1/2}).$$

Furthermore, the asymptotic distribution of $\hat{\theta}_{\text{SUP}}$ is:

$$n^{1/2} f(\theta_0) \sigma_{\text{SUP}}^{-1} (\hat{\theta}_{\text{SUP}} - \theta_0) \rightarrow Normal(0, 1) \quad (n \rightarrow \infty),$$

with $\sigma_{\text{SUP}}^2 := \text{var}\{\psi(Y, \theta_0)\} = \tau(1 - \tau)$.

Proposition 2.1 provides the asymptotic variance of $\hat{\theta}_{\text{SUP}}$, which can be compared with that of our semi-supervised estimator(s) proposed below. Its proof can be found in [Koenker \(2005\)](#).

2.2 A family of semi-supervised estimators based on one-step update

Main idea. The conditional expectation of the left hand side in (1.1) given \mathbf{X} is:

$$\mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\} \equiv F(\theta_0 \mid \mathbf{X}) - \tau \neq 0 \text{ with a positive probability,}$$

if we exclude the trivial situation where $F(\theta_0 \mid \mathbf{X}) = \tau$ almost surely. This indicates that the distribution $\mathbb{P}_{\mathbf{X}}$ of \mathbf{X} indeed plays a role in the definition of θ_0 . Hence, the supervised estimator $\hat{\theta}_{\text{SUP}}$ is possibly sub-optimal in that it discards the unlabeled data \mathcal{U} , which provides extra information on $\mathbb{P}_{\mathbf{X}}$. To make use of \mathcal{U} where Y is not observed, we now consider substituting a function of \mathbf{X} for $\mathbb{E}\{\psi(Y, \theta)\}$ in the definition (1.1) of θ_0 . An intuitive choice is $\mu(\mathbf{X}, \theta) := \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$ since $\mathbb{E}\{\psi(Y, \theta)\} = \mathbb{E}\{\mu(\mathbf{X}, \theta)\}$. Then, a further representation of $\mathbb{E}\{\psi(Y, \theta)\}$ is:

$$\mathbb{E}\{\psi(Y, \theta)\} = \mathbb{E}\{\mu(\mathbf{X}, \theta)\} + \mathbb{E}\{\psi(Y, \theta) - \mu(\mathbf{X}, \theta)\}.$$

However, the form of $\mu(\mathbf{X}, \theta)$ is typically hard to specify correctly in practice. We thus posit a *working model* $\phi(\mathbf{X}, \theta)$ that is possibly misspecified, and *not* necessarily equal to $\mu(\mathbf{X}, \theta)$.

Then, our method is inspired by the fact that the following *robust representation* holds:

$$\mathbb{E}\{\psi(Y, \theta)\} = h(\theta) := \mathbb{E}\{\phi(\mathbf{X}, \theta)\} + \mathbb{E}\{\psi(Y, \theta) - \phi(\mathbf{X}, \theta)\}, \quad (2.1)$$

for an *arbitrary* function $\phi(\cdot, \cdot)$, implying that under (1.1), at $\theta = \theta_0$,

$$h(\theta_0) \equiv \mathbb{E}\{\phi(\mathbf{X}, \theta_0)\} + \mathbb{E}\{\psi(Y, \theta_0) - \phi(\mathbf{X}, \theta_0)\} = 0. \quad (2.2)$$

In the left hand side of (2.2), the term $\mathbb{E}\{\psi(Y, \theta_0) - \phi(\mathbf{X}, \theta_0)\}$ guarantees *global robustness* which means the equation always holds, i.e., for any function $\phi(\cdot, \cdot)$, while the other term $\mathbb{E}\{\phi(\mathbf{X}, \theta_0)\}$ involves \mathbf{X} only and can thus be estimated using the whole data $\mathcal{L} \cup \mathcal{U}$.

Then, the sample version of (2.2) constructed with $\mathcal{L} \cup \mathcal{U}$ is:

$$\mathbb{E}_{n+N}\{\hat{\phi}_n(\mathbf{X}, \theta)\} + \mathbb{E}_n\{\psi(Y, \theta) - \hat{\phi}_n(\mathbf{X}, \theta)\} = 0, \quad (2.3)$$

where $\hat{\phi}_n(\cdot, \cdot)$ denotes some reasonable estimator of $\phi(\cdot, \cdot)$ based on \mathcal{L} .

Construction of the estimators. Intuitively, the equation (2.3) gives the road map to our semi-supervised estimators. Solving (2.3) with respect to θ is, however, not straightforward owing to its non-linear nature. So we consider a more implementation-friendly and computationally efficient one-step update approach. Noticing that the derivative $h'(\cdot)$ of the function $h(\cdot)$ defined in (2.1) is the density function $f(\cdot)$ of Y , we can hence solve the equation (2.2) by Newton's method, which refines an initial solution θ_{INIT} by a one-step update $\theta_{\text{INIT}} - \{h'(\theta_{\text{INIT}})\}^{-1}h(\theta_{\text{INIT}}) \equiv$

$$\theta_{\text{INIT}} + \{f(\theta_{\text{INIT}})\}^{-1}[\mathbb{E}\{\phi(\mathbf{X}, \theta_{\text{INIT}}) - \psi(Y, \theta_{\text{INIT}})\} - \mathbb{E}\{\phi(\mathbf{X}, \theta_{\text{INIT}})\}]. \quad (2.4)$$

Recall from (2.3) $\hat{\phi}_n(\cdot, \cdot)$ is an estimator of $\phi(\cdot, \cdot)$ based on \mathcal{L} . Further, let $\hat{\theta}_{\text{INIT}}$ be an initial estimator of θ_0 and let $\hat{f}_n(\cdot)$ be an estimator of $f(\cdot)$. Then, the empirical version, based on the whole data

$\mathcal{L} \cup \mathcal{U}$, of the population-level representation (2.4) immediately gives a family of *semi-supervised estimators* $\hat{\theta}_{\text{ss}}$ of θ_0 indexed by $\{\hat{\phi}_n(\cdot, \cdot), \hat{\theta}_{\text{INIT}}, \hat{f}_n(\cdot)\}$:

$$\hat{\theta}_{\text{ss}} := \hat{\theta}_{\text{INIT}} + \{\hat{f}_n(\hat{\theta}_{\text{INIT}})\}^{-1} [\mathbb{E}_n\{\hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}}) - \psi(Y, \hat{\theta}_{\text{INIT}})\} - \mathbb{E}_{n+N}\{\hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}})\}]. \quad (2.5)$$

Although we do not require specific forms of $\{\hat{\theta}_{\text{INIT}}, \hat{f}_n(\cdot)\}$, a natural choice of the initial estimator $\hat{\theta}_{\text{INIT}}$ is the supervised estimator $\hat{\theta}_{\text{SUP}}$, while $\hat{f}_n(\cdot)$ can be a kernel density estimator; see Remark 3.2 for details on their convergence properties. As regards the imputation function $\hat{\phi}_n(\cdot, \cdot)$, which is an important component in our method, an arbitrary choice is allowed as long as the high-level conditions in Section 2.3 are satisfied. We will thoroughly study some specific examples of $\hat{\phi}_n(\cdot, \cdot)$ in Section 3. However, regardless of the choice of $\hat{\phi}_n(\cdot, \cdot)$, we apply a general *cross-fitting* strategy (Chernozhukov et al., 2018; Newey and Robins, 2018) to obtain $\hat{\phi}_n(\mathbf{X}_i, \cdot)$ as follows.

Cross-fitting of $\hat{\phi}_n(\cdot, \cdot)$ and its benefits. For some fixed integer $\mathbb{K} \geq 2$, we divide the index set $\mathcal{I} = \{1, \dots, n\}$ into \mathbb{K} disjoint subsets $\mathcal{I}_1, \dots, \mathcal{I}_{\mathbb{K}}$ of the same size $n_{\mathbb{K}} := n/\mathbb{K}$ without loss of generality. Let $\hat{\phi}_{n,k}(\cdot, \cdot)$ be the corresponding estimator of $\phi(\cdot, \cdot)$ based on the data $\mathcal{L}_k^- := \{\mathbf{Z}_i : i \in \mathcal{I}_k^-\}$ of size $n_{\mathbb{K}^-} := n - n_{\mathbb{K}}$, where $\mathcal{I}_k^- := \mathcal{I}/\mathcal{I}_k$. Then we set

$$\hat{\phi}_n(\mathbf{X}_i, \theta) \equiv \sum_{k=1}^{\mathbb{K}} \{\hat{\phi}_{n,k}(\mathbf{X}_i, \theta) I(i \in \mathcal{I}_k) + \mathbb{K}^{-1} \hat{\phi}_{n,k}(\mathbf{X}_i, \theta) I(i > n)\} \quad (2.6)$$

Through cross-fitting, the dependence of $\hat{\phi}_n(\cdot, \cdot)$ and \mathbf{X}_i in $\hat{\phi}_n(\mathbf{X}_i, \theta)$ ($i = 1, \dots, n$) is eliminated, so that the second-order errors in the expansion of $\hat{\theta}_{\text{ss}}$ become more tractable while the influence function remains unchanged. We can therefore avoid some stringent conditions, which are analogous to the stochastic equicontinuity ones in the empirical process theory (Van der Vaart, 2000), when deriving properties of $\hat{\theta}_{\text{ss}}$. More detailed discussions of cross-fitting can be found in Chakrabortty and Cai (2018), Chernozhukov et al. (2018) and Newey and Robins (2018).

In summary, we obtain our semi-supervised estimators $\hat{\theta}_{\text{ss}}$ of the quantile θ_0 in three steps:

- (i) Calculate an initial estimator $\hat{\theta}_{\text{INIT}}$ of θ_0 and an estimator $\hat{f}_n(\cdot)$ of the density function $f(\cdot)$;
- (ii) Obtain the imputation function $\hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}})$ by the cross-fitting procedures (2.6);
- (iii) Plug $\{\hat{f}_n(\hat{\theta}_{\text{INIT}}), \hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}})\}$ in the one-step update formula (2.5) to obtain the final $\hat{\theta}_{\text{ss}}$.

Remark 2.1 (Robustifying and debiasing nature of the representation (2.5)). In addition to robustifying the estimator $\hat{\theta}_{\text{ss}}$, as discussed after (2.2), another effect of the term $\mathbb{E}_n\{\hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}}) - \psi(Y, \hat{\theta}_{\text{INIT}})\}$ in (2.5) is eradicating the first-order error of $\hat{\phi}_n(\cdot, \cdot)$ as an estimator of $\phi(\cdot, \cdot)$ so that the influence function of $\hat{\theta}_{\text{ss}}$ is not affected by the estimation error of $\hat{\phi}_n(\cdot, \cdot)$. This property is crucial for ensuring the $n^{1/2}$ -consistency and asymptotic normality of $\hat{\theta}_{\text{ss}}$, particularly when \mathbf{X} is high dimensional or $\hat{\phi}_n(\cdot, \cdot)$ involves nonparametric calibrations. We will formally discuss this point in Remark 2.4 after obtaining the theoretical results of $\hat{\theta}_{\text{ss}}$. Interestingly, even if the imputation function satisfies $\mathbb{E}\{\phi(\mathbf{X}, \theta)\} = \mathbb{E}\{\psi(Y, \theta)\}$, this *robustifying and debiasing* term should *always* be included so that $\hat{\theta}_{\text{ss}}$ can enjoy the desirable properties mentioned above.

2.3 Main results: Theoretical properties of the semi-supervised estimators

The definition (2.5) now equips us with a family of semi-supervised estimators $\widehat{\theta}_{\text{ss}}$ for θ_0 indexed by $\{\widehat{\phi}_n(\cdot, \cdot), \widehat{\theta}_{\text{INIT}}, \widehat{f}_n(\cdot)\}$. To study their limiting behavior, we assume the following conditions.

Assumption 2.2. The estimators $\widehat{\theta}_{\text{INIT}}$ and $\widehat{f}_n(\cdot)$ satisfy that

$$\widehat{\theta}_{\text{INIT}} - \theta_0 = O_p(u_n) \text{ and } \widehat{f}_n(\widehat{\theta}_{\text{INIT}}) - f(\theta_0) = O_p(v_n)$$

for some positive sequences $u_n = o(1)$ and $v_n = o(1)$.

Assumption 2.3. $\mathbb{E}[\{\phi(\mathbf{X}, \theta_0)\}^2] < \infty$. In addition, for any sequence $\widetilde{\theta} \rightarrow \theta_0$ in probability,

$$\mathbb{G}_n\{\phi(\mathbf{X}, \widetilde{\theta}) - \phi(\mathbf{X}, \theta_0)\} = o_p(1) \text{ and } \mathbb{G}_{n+N}\{\phi(\mathbf{X}, \widetilde{\theta}) - \phi(\mathbf{X}, \theta_0)\} = o_p(1). \quad (2.7)$$

Assumption 2.4. Denote the estimation error $\widehat{\psi}_{n,k}(\mathbf{X}, \theta) := \widehat{\phi}_{n,k}(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta)$ and its second moment $\Delta_k(\mathcal{L}) := (\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} \mathbb{E}_{\mathbf{X}}[\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}^2])^{1/2}$ ($k = 1, \dots, \mathbb{K}$). Then, the set

$$\mathcal{P}_{n,k} := \{\widehat{\psi}_{n,k}(\mathbf{X}, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}, \quad (2.8)$$

for some $\varepsilon > 0$, satisfies that, for any $\eta \in (0, \Delta_k(\mathcal{L}) + \xi]$ with some $\xi > 0$,

$$N_{[]} \{\eta, \mathcal{P}_{n,k} \mid \mathcal{L}, L_2(\mathbb{P}_{\mathbf{X}})\} \leq H(\mathcal{L})\eta^{-c} \quad (k = 1, \dots, \mathbb{K}), \quad (2.9)$$

with some function $H(\mathcal{L}) > 0$ such that $H(\mathcal{L}) = O_p(a_n)$ for some positive sequence a_n , where the symbol $N_{[]}(\cdot, \cdot, \cdot)$ refers to the bracketing number defined in [Van der Vaart and Wellner \(1996\)](#) and [Van der Vaart \(2000\)](#). Here $\mathcal{P}_{n,k}$ is indexed by θ only and treats $\widehat{\psi}_{n,k}(\cdot, \theta)$ as a non-random function ($k = 1, \dots, \mathbb{K}$). Further, for some positive sequences $d_{n,2}$ and $d_{n,\infty}$ allowed to diverge,

$$\Delta_k(\mathcal{L}) = O_p(d_{n,2}) \text{ and } \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{\psi}_{n,k}(\mathbf{x}, \theta)| = O_p(d_{n,\infty}) \quad (k = 1, \dots, \mathbb{K}). \quad (2.10)$$

Remark 2.2. Assumption 2.2 is standard for one-step estimators, ensuring good behavior of $\widehat{\theta}_{\text{INIT}}$ and $\widehat{f}_n(\cdot)$. Assumption 2.3 outlines features of a reasonable imputation function $\phi(\cdot, \cdot)$. According to Example 19.7 and Lemma 19.24 of [Van der Vaart \(2000\)](#), the condition (2.7) is true provided $\phi(\mathbf{X}, \theta)$ is Lipschitz continuous in θ . Assumption 2.4 is imposed to control the estimation error of $\widehat{\phi}_{n,k}(\mathbf{X}, \theta)$ in the neighborhood $\mathcal{B}(\theta_0, \varepsilon)$ of θ_0 . The condition (2.9) therein holds when $\mathcal{P}_{n,k}$ is a VC class given \mathcal{L} ([Van der Vaart and Wellner, 1996](#)). Also, we put the restriction (2.10) with possibly divergent rates on the L_2 and L_∞ norms of $\widehat{\psi}_{n,k}(\mathbf{X}, \theta)$, *weaker* than requiring its convergence uniformly over $\mathbf{x} \in \mathcal{X}$ and $\theta \in \mathcal{B}(\theta_0, \varepsilon)$, i.e., the L_∞ convergence. All these (high-level) assumptions are fairly mild and will be verified for some choices of $\{\phi(\cdot, \cdot), \widehat{\phi}_{n,k}(\cdot, \cdot)\}$ in Section 3; see, e.g., [Propositions 3.1–3.2](#) and the discussions after [Theorem 3.1](#) therein. In addition, we do *not* assume $\lim_{n \rightarrow \infty} (n/N) = 0$, a common requirement in the semi-supervised literature ([Chakrabortty and Cai, 2018; Gronsbell and Cai, 2018](#)). Our conclusions thus remain valid even when the labeled and unlabeled data are comparable in size. Nevertheless $\nu = 0$ is a more practically relevant and theoretically challenging case, significantly different from the traditional missing data problem.

In the following theorem, we state the large sample properties of $\widehat{\theta}_{\text{ss}}$ defined by (2.5), giving a complete characterization of its asymptotic expansion under general (high-level) conditions.

Theorem 2.1 (General asymptotic expansion of $\hat{\theta}_{ss}$). *If Assumptions 2.1–2.4 hold, then*

$$\hat{\theta}_{ss} - \theta_0 = \{nf(\theta_0)\}^{-1} \sum_{i=1}^n \omega_{n,N}(\mathbf{Z}_i, \theta_0) + O_p(u_n^2 + u_n v_n + n^{-1/2} r_n) + o_p(n^{-1/2}),$$

where $\omega_{n,N}(\mathbf{Z}, \theta) := \phi(\mathbf{X}, \theta) - \psi(Y, \theta) - \mathbb{E}_{n+N}\{\phi(\mathbf{X}, \theta)\}$ satisfying $\mathbb{E}\{\omega_{n,N}(\mathbf{Z}, \theta)\} = 0$, and

$$r_n := d_{n,2}\{\log a_n + \log(d_{n,2}^{-1})\} + n_{\mathbb{K}}^{-1/2} d_{n,\infty}\{(\log a_n)^2 + (\log d_{n,2})^2\}.$$

Further, given

$$u_n v_n + u_n^2 + n^{-1/2} r_n = o(n^{-1/2}), \quad (2.11)$$

the limiting distribution of $\hat{\theta}_{ss}$ is $n^{1/2} f(\theta_0) \sigma_{ss}^{-1}(\hat{\theta}_{ss} - \theta_0) \rightarrow \text{Normal}(0, 1)$ ($n, N \rightarrow \infty$), where

$$\sigma_{ss}^2 := \text{var}\{\omega_{n,N}(\mathbf{Z}, \theta_0)\} = (1 - \nu_{n,N}) \text{var}\{\psi(Y, \theta_0) - \phi(\mathbf{X}, \theta_0)\} + \nu_{n,N} \text{var}\{\psi(Y, \theta_0)\}.$$

Remark 2.3. The asymptotic variance of $\hat{\theta}_{ss}$ can be estimated by $n^{-1}\{\hat{f}_n(\hat{\theta}_{\text{INIT}})\}^{-2}\hat{\sigma}_{ss}^2$ with

$$\hat{\sigma}_{ss}^2 := (1 - \nu_{n,N}) \text{var}_n\{\psi(Y, \hat{\theta}_{\text{INIT}}) - \hat{\phi}_n(\mathbf{X}, \hat{\theta}_{\text{INIT}})\} + \nu_{n,N} \text{var}_n\{\psi(Y, \hat{\theta}_{\text{INIT}})\}. \quad (2.12)$$

Of course, one can replace the initial estimator $\hat{\theta}_{\text{INIT}}$ in $\hat{f}_n(\hat{\theta}_{\text{INIT}})$ and (2.12) by the semi-supervised estimator $\hat{\theta}_{ss}$. The results from our simulations in Section 4, however, show that $\{\hat{f}_n(\hat{\theta}_{\text{INIT}})\}^{-1}\hat{\sigma}_{ss}$ works quite well for estimating $\{f(\theta_0)\}^{-1}\sigma_{ss}$. In addition, using $\{\hat{f}_n(\hat{\theta}_{\text{INIT}})\}^{-1}\hat{\sigma}_{ss}$ reduces computational burden, since $\hat{\phi}_n(\mathbf{X}_i, \hat{\theta}_{\text{INIT}})$ is already available from the previous steps, while $\hat{\phi}_n(\mathbf{X}_i, \hat{\theta}_{ss})$ ($i = 1, \dots, n$) needs to be calculated afresh via the cross-fitting procedure (2.6).

As the most important special case of Theorem 2.1, which holds for any $\nu_{n,N}$ and its limit ν – positive or zero, the limiting behavior of $\hat{\theta}_{ss}$ when $\nu = 0$ is considered in the next corollary.

Corollary 2.1 (Properties of $\hat{\theta}_{ss}$ under the special case $\nu = 0$). *Assume that the conditions of Theorem 2.1 hold and that $\nu = 0$. Then, the semi-supervised estimator $\hat{\theta}_{ss}$ satisfies:*

$$\begin{aligned} \hat{\theta}_{ss} - \theta_0 &= \{nf(\theta_0)\}^{-1} \sum_{i=1}^n \omega(\mathbf{Z}_i, \theta_0) + O_p(u_n^2 + u_n v_n + n^{-1/2} r_n) + o_p(n^{-1/2}), \text{ and} \\ n^{1/2} f(\theta_0) \tilde{\sigma}_{ss}^{-1}(\hat{\theta}_{ss} - \theta_0) &\rightarrow \text{Normal}(0, 1) \quad (n, N \rightarrow \infty), \end{aligned}$$

where $\omega(\mathbf{Z}, \theta) := \phi(\mathbf{X}, \theta) - \psi(Y, \theta) - \mathbb{E}\{\phi(\mathbf{X}, \theta)\}$ satisfying $\mathbb{E}\{\omega(\mathbf{Z}, \theta)\} = 0$, and

$$\tilde{\sigma}_{ss}^2 := \text{var}[\{\omega(\mathbf{Z}, \theta_0)\}^2] = \text{var}\{\psi(Y, \theta_0) - \phi(\mathbf{X}, \theta_0)\}.$$

Remark 2.4 (Robustness and first-order insensitivity). Theorem 2.1 presents the $n^{1/2}$ -consistency and asymptotic normality of $\hat{\theta}_{ss}$ with an *arbitrary* choice of $\{\phi(\cdot, \cdot), \hat{\phi}_{n,k}(\cdot, \cdot)\}$ under the assumptions therein. In this sense, it provides a *family of globally robust semi-supervised estimators* with influence functions indexed by $\phi(\cdot, \cdot)$. In addition, we observe that estimating $\phi(\cdot, \cdot)$ by $\hat{\phi}_n(\cdot, \cdot)$ does *not* affect the influence function of $\hat{\theta}_{ss}$, as long as the high-level conditions in Assumption 2.4 are satisfied. Therefore, $\hat{\theta}_{ss}$ is first-order insensitive to estimation errors and any knowledge of the construction of $\hat{\phi}_n(\cdot, \cdot)$. This property is particularly desirable for inference when \mathbf{X} is high dimensional or $\hat{\phi}_n(\cdot, \cdot)$ involves nonparametric techniques – cases when it may not be $n^{-1/2}$ -rate.

Remark 2.5 (Efficiency comparisons, and some examples of the imputation function $\phi(\cdot, \cdot)$). If we take $\phi(\mathbf{X}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{d}(\mathbf{X})\}$ with some possibly unknown function $\mathbf{d}(\cdot)$, then

$$\sigma_{\text{ss}}^2 \equiv \mathbb{E}[\{\psi(Y, \theta_0)\}^2] - (1 - \nu_{n,N})\mathbb{E}[\{\phi(\mathbf{X}, \theta_0)\}^2] \leq \sigma_{\text{sup}}^2,$$

i.e., the semi-supervised variance $\{f(\theta_0)\}^{-2}\sigma_{\text{ss}}^2$ in Theorem 2.1 is no more than the supervised variance $\{f(\theta_0)\}^{-2}\sigma_{\text{sup}}^2$ in Proposition 2.1, indicating $\hat{\theta}_{\text{ss}}$ is *equally or more efficient* compared to the supervised estimator $\hat{\theta}_{\text{sup}}$. An example of $\mathbf{d}(\mathbf{x})$ is the linear transformation $\mathbf{d}(\mathbf{x}) \equiv \mathbf{P}_0^T \mathbf{x}$, where \mathbf{P}_0 is some unknown $r \times p$ matrix, with a fixed $r \leq p$, that can be chosen and estimated using parametric regression methods ($r = 1$), e.g., linear regression of Y vs. \mathbf{X} , or dimension reduction techniques ($r \geq 1$) such as sliced inverse regression (Li, 1991; Lin et al., 2019); see Section 4 for the implementation details of estimating \mathbf{P}_0 . After obtaining an estimator of \mathbf{P}_0 , a further step of nonparametric smoothing can be conducted to approximate $\phi(\mathbf{x}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{P}_0^T \mathbf{X} = \mathbf{P}_0^T \mathbf{x}\}$. In Section 3, we will substantiate that $\phi(\mathbf{x}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{P}_0^T \mathbf{X} = \mathbf{P}_0^T \mathbf{x}\}$ and its corresponding nuisance estimators indeed satisfy the high-level conditions required in Theorem 2.1.

Further, when $\phi(\mathbf{X}, \theta_0) = \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}$ and $\nu > 0$, we have:

$$\begin{aligned} \sigma_{\text{ss}}^2 &= (1 - \nu_{n,N})\text{var}[\psi(Y, \theta_0) - \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}] + \nu_{n,N}\text{var}\{\psi(Y, \theta_0)\} \\ &\rightarrow (1 - \nu)\mathbb{E}([\psi(Y, \theta_0) - \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}]^2) + \nu\mathbb{E}[\{\psi(Y, \theta_0)\}^2] = \sigma_{\text{eff}}^2, \end{aligned} \quad (2.13)$$

with $\{f(\theta_0)\}^{-2}\sigma_{\text{eff}}^2$ the *semi-parametric efficiency bound* for estimating θ_0 in missing data theory (Tsiatis, 2007; Graham, 2011). If $\phi(\mathbf{X}, \theta_0) = \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}$ and $\nu = 0$, Corollary 2.1 gives:

$$\tilde{\sigma}_{\text{ss}}^2 = \mathbb{E}([\psi(Y, \theta_0) - \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}]^2) \leq \mathbb{E}[\{\psi(Y, \theta_0) - g(\mathbf{X})\}^2],$$

for any function $g(\cdot)$ in $L_2(\mathbb{P}_{\mathbf{X}})$, and the equality holds only if $g(\mathbf{X}) = \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}$ almost surely. This fact reveals the asymptotic *optimality* of $\hat{\theta}_{\text{ss}}$ among all regular and asymptotically linear estimators of θ_0 , whose influence functions take the form: $\{f(\theta_0)\}^{-1}\{g(\mathbf{X}) - \psi(Y, \theta_0)\}$, for some function $g(\cdot)$. Further, under the appropriate semi-parametric model of \mathbf{Z} – one where the distribution of \mathbf{X} is known while that of Y is unrestricted up to Assumption 2.1, one can show that $\{f(\theta_0)\}^{-2}\mathbb{E}([\psi(Y, \theta_0) - \mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{X}\}]^2)$ equals the *efficient* asymptotic variance for estimating θ_0 . Thus, for any $\nu \geq 0$, $\hat{\theta}_{\text{ss}}$ achieves *semi-parametric efficiency* with $\phi(\cdot, \cdot)$ as above.

3 Choice and estimation of the nuisance functions

This section details some choices and estimators of the imputation function $\phi(\cdot, \cdot)$ used in the construction of our semi-supervised estimators $\hat{\theta}_{\text{ss}}$ in (2.5). Although an arbitrary imputation function equips our method with the $n^{1/2}$ -consistency and asymptotic normality stated in Theorem 2.1 if the high-level conditions in Assumption 2.3 hold, the ideal choice from the perspective of efficiency is $\phi(\mathbf{X}, \theta) = \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$ as discussed in Remark 2.5. However, when the dimension p of \mathbf{X} is large, estimating the conditional mean $\mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$ fully nonparametrically is generally undesirable in practice due to curse of dimensionality that typically enforces stringent conditions such as undersmoothing (Chakrabortty and Cai, 2018). A common strategy is implementing suitable *dimension reduction* techniques followed by *nonparametric calibrations* targeting the function $\mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S}\}$ rather than $\mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$, where $\mathbf{S} := \mathbf{P}_0^T \mathbf{X} \in \mathcal{S} \subset \mathbb{R}^r$ and \mathbf{P}_0 is a $p \times r$ matrix with some fixed $r \leq p$. Here we emphasize that $\mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S}\} = \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$ is *not* assumed anyway, i.e., the

dimension reduction is *not* necessarily sufficient. According to Remark 2.5, the advantage of $\widehat{\theta}_{\text{ss}}$ over $\widehat{\theta}_{\text{SUP}}$ in terms of efficiency is ensured by setting

$$\phi(\mathbf{X}, \theta) \equiv \phi(\mathbf{X}, \theta, \mathbf{P}_0) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{P}_0^T \mathbf{X}\} \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S}\}, \quad (3.1)$$

regardless of whether the dimension reduction is sufficient or not. Hence \mathbf{P}_0 can be *any* user-defined or data-dependent matrix. If \mathbf{P}_0 is entirely determined by $\mathbb{P}_{\mathbf{X}}$, we can safely assume its estimation error to be negligible due to the plentiful observations for \mathbf{X} in semi-supervised settings. An example is the r leading principal component directions of \mathbf{X} . However, to make the dimension reduction and the imputation as sufficient as possible, we mainly consider cases where \mathbf{P}_0 *depends* on the joint distribution of $(Y, \mathbf{X}^T)^T$ and thereby needs to be estimated on \mathcal{L} with significant errors. Some reasonable choices of such \mathbf{P}_0 will be discussed in Remark 3.1.

To justify the usage of imputation functions of the form $\phi(\mathbf{X}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S}\}$ as in (3.1), we now show it satisfies Assumption 2.3 under a mild condition.

Proposition 3.1. *Suppose that $\mathbb{E}[\{\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} f(\theta \mid \mathbf{S})\}^2] < \infty$ for some ε , with $f(\cdot \mid \mathbf{S})$ the conditional density of Y given \mathbf{S} . Then, Assumption 2.3 is satisfied by $\phi(\mathbf{X}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S}\}$.*

To approximate the conditional mean $\phi(\mathbf{X}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{P}_0^T \mathbf{X}\}$, we may employ any suitable smoothing technique, such as kernel smoothing, kernel machine regression and smoothing splines. For sake of illustration, we focus here on the *kernel smoothing estimator(s)*:

$$\begin{aligned} \widehat{\phi}_{n,k}(\mathbf{x}, \theta) &\equiv \widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) := \{\widehat{\ell}_{n,k}(\mathbf{x}, \widehat{\mathbf{P}}_k)\}^{-1} \widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) \quad (k = 1, \dots, \mathbb{K}), \text{ where} \\ \widehat{\ell}_{n,k}(\mathbf{x}, \mathbf{P}) &:= h_n^{-r} \mathbb{E}_{n,k}[K_h\{\mathbf{P}^T(\mathbf{x} - \mathbf{X})\}] \text{ and } \widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}) := h_n^{-r} \mathbb{E}_{n,k}[\psi(Y, \theta) K_h\{\mathbf{P}^T(\mathbf{x} - \mathbf{X})\}], \end{aligned} \quad (3.2)$$

with the notation $\mathbb{E}_{n,k}\{g(\mathbf{Z})\} := n_{\mathbb{K}}^{-1} \sum_{i \in \mathcal{I}_k^-} g(\mathbf{Z}_i)$ for any function $g(\cdot)$, and with $\widehat{\mathbf{P}}_k$ being any estimator of \mathbf{P}_0 based on the data set \mathcal{L}_k^- , $K_h(\mathbf{s}) := K(h_n^{-1} \mathbf{s})$, $K(\cdot) : \mathbb{R}^r \mapsto \mathbb{R}$ a kernel function, e.g., the standard Gaussian kernel, and $h_n \rightarrow 0$ denoting a bandwidth sequence. Considering \mathbf{X} is possibly high dimensional and \mathbf{P}_0 needs to be estimated, establishing the (uniform) convergence properties of $\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k)$ poses substantial technical challenges and has not been studied in the literature yet to the best of our knowledge. In contrast, most of the existing works on kernel smoothing with such estimated transformed covariates – often termed as “generated” covariates – e.g., Mammen et al. (2012), Escanciano et al. (2014) and Mammen et al. (2016), mainly focus on the scenario where the dimension of \mathbf{X} is fixed. Our result below is thus *novel* in this sense.

To derive the convergence rate of $\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k)$, as defined in (3.2), *uniformly* over $\mathbf{x} \in \mathcal{X}$ and $\theta \in \mathcal{B}(\theta_0, \varepsilon)$, we impose the following requirements, along with some standard smoothness conditions for kernel smoothing which are listed in Assumption A.1 of Appendix A.1.

Assumption 3.1. The estimation error $\|\widehat{\mathbf{P}}_k - \mathbf{P}_0\|_1 = O_p(\alpha_n)$ for some positive sequence α_n .

Assumption 3.2 (Required only when \mathbf{P}_0 needs to be estimated). (i) The support \mathcal{X} of \mathbf{X} is such that $\sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_\infty < \infty$. (ii) The function $\nabla K(\mathbf{s}) := \partial K(\mathbf{s}) / \partial \mathbf{s}$ is continuously differentiable and satisfies $\|\partial\{\nabla K(\mathbf{s})\} / \partial \mathbf{s}\| \leq c \|\mathbf{s}\|^{-w}$ for any $\|\mathbf{s}\| > c$, where $w > 1$ is some constant. Further, it is locally lipschitz continuous, i.e., $\|\nabla K(\mathbf{s}_1) - \nabla K(\mathbf{s}_2)\| \leq \|\mathbf{s}_1 - \mathbf{s}_2\| \rho(\mathbf{s}_2)$ for any $\|\mathbf{s}_1 - \mathbf{s}_2\| \leq c$, where $\rho(\cdot)$ is some bounded and square integrable function with a bounded derivative $\nabla \rho(\cdot)$. (iii) Let $\boldsymbol{\eta}_{t[j]}(\mathbf{s}, \theta)$ be the j th component of $\boldsymbol{\eta}_t(\mathbf{s}, \theta) := \mathbb{E}[\mathbf{X}\{\psi(Y, \theta)\}^t \mid \mathbf{S} = \mathbf{s}]$. Then, with respect to \mathbf{s} , the function $\boldsymbol{\eta}_{t[j]}(\mathbf{s}, \theta)$ is continuously differentiable and has a bounded first derivative on $\mathcal{S}_0 \times \mathcal{B}(\theta_0, \varepsilon)$ ($t = 0, 1; j = 1, \dots, p$) for some open set $\mathcal{S}_0 \supset \mathcal{S}$.

Assumption 3.1 regulates the behavior of $\widehat{\mathbf{P}}_k$ as an estimator of the transformation matrix \mathbf{P}_0 . Assumption 3.2 requires mild smoothness conditions to control the estimation error of $\widehat{\mathbf{P}}_k$ while (ii) therein is satisfied in particular by the second-order Gaussian kernel, among others. Similar assumptions can be found in Chakrabortty and Cai (2018) that studied kernel smoothing estimators with dimension reduction in low dimensional scenarios. We now propose the following result.

Theorem 3.1 (Convergence of $\widehat{\phi}_{n,k}$). *Set $\gamma_n := [(nh_n^r)^{-1} \max\{\log(h_n^{-r}), \log(\log n)\}]^{1/2}$, $s_{n,1} := \gamma_n + h_n^d$ and $s_{n,2} := h_n^{-2}\alpha_n^2 + h_n^{-1}\gamma_n\alpha_n + \alpha_n$. If Assumptions 3.1–3.2, as well as Assumption A.1 in Appendix A.1, hold and $s_{n,1} + s_{n,2} = o(1)$, then*

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - \phi(\mathbf{x}, \theta, \mathbf{P}_0)| = O_p(s_{n,1} + s_{n,2}) \quad (k = 1, \dots, \mathbb{K}).$$

Remark 3.1 (Convergence rates – examples of $\widehat{\mathbf{P}}_k$). Theorem 3.1 establishes the L_∞ error rate of $\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k)$ under mild conditions. The uniform consistency of $\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k)$ is ensured at the optimal bandwidth rate $h_{\text{opt}} = O\{n^{-1/(2d+r)}\}$ for any kernel order $d \geq 2$ and fixed r , if

$$\alpha_n = o\{n^{-1/(2d+r)}\}. \quad (3.3)$$

We consider the validity of (3.3) for some common choices of \mathbf{P}_0 , such as the least squares regression parameter ($r = 1$) satisfying $\mathbb{E}\{\mathbf{X}(Y - \mathbf{P}_0^T \mathbf{X})\} = \mathbf{0}_p$, or the r leading eigenvectors ($r \geq 1$) of the matrix $\text{cov}\{\mathbb{E}(\mathbf{X} | Y)\}$, which can be estimated by sliced inverse regression (Li, 1991). When p is fixed, there typically exist $n^{1/2}$ -consistent estimators $\widehat{\mathbf{P}}_k$ of \mathbf{P}_0 , so (3.3) is satisfied by the fact that $\alpha_n = O(n^{-1/2})$. In high dimensional scenarios where p is divergent and greater than n , one can obtain $\widehat{\mathbf{P}}_k$ based on regularized versions of linear regression or sliced inverse regression (Lin et al., 2019). In these cases, the sequence $\alpha_n = O\{q(\log p/n)^{1/2}\}$ when the L_1 penalty is used under some suitable conditions (Bühlmann and Van De Geer, 2011; Wainwright, 2019; Lin et al., 2019), where $q := \|\mathbf{P}_0\|_0$ represents the sparsity level. Hence (3.3) is true whenever $q(\log p)^{1/2} = o\{n^{(2d+r-2)/(4d+2r)}\}$. An example of $\widehat{\mathbf{P}}_k$, with $r = 1$, for high dimensional data is the lasso estimator $\widehat{\mathbf{P}}_k \equiv \arg \min_{\beta \in \mathbb{R}^p} \{n_{\mathbb{K}^-}^{-1} \sum_{i \in \mathcal{I}_k^-} (Y_i - \beta^T \mathbf{X}_i)^2 + \lambda_{n,k} \|\beta\|_1\}$, where $\lambda_{n,k} > 0$ is a tuning parameter. After this parametric regression step, we conduct kernel smoothing of $\psi(Y, \widehat{\theta}_{\text{INIT}})$ on the one-dimensional linear combination $\widehat{\mathbf{P}}_k^T \mathbf{X}$ to obtain $\widehat{\phi}_{n,k}(\mathbf{x}, \widehat{\theta}_{\text{INIT}}, \widehat{\mathbf{P}}_k)$, as in (3.2); see Section 4 for the details of constructing such a high-dimensional nuisance estimator.

Theorem 3.1 implies that $d_{n,2} = o(1)$ and $d_{n,\infty} = o(1)$ for Assumption 2.4 when $\phi(\mathbf{x}, \theta, \mathbf{P})$ and $\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \mathbf{P})$ are as in (3.1)–(3.2). We now validate the condition (2.9) on the bracketing number.

Proposition 3.2. *Suppose that $\mathbb{E}[\{\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} f(\theta | \mathbf{S})\}^2] < \infty$, for some $\varepsilon > 0$. Then, the set $\mathcal{P}_{n,k}$ defined in (2.8) satisfies $N_{[]} \{\eta, \mathcal{P}_{n,k} | \mathcal{L}, L_2(\mathbb{P}_{\mathbf{X}})\} \leq c(n+1)\eta^{-1}$.*

Remark 3.2 (Verification of the condition (2.11)). The results of Theorem 3.1 and Proposition 3.2 indicate that $a_n = O(n)$, $d_{n,2} = o(1)$ and $d_{n,\infty} = o(1)$ in Assumption 2.4. Moreover, by setting the initial estimator $\widehat{\theta}_{\text{INIT}} \equiv \widehat{\theta}_{\text{SUP}}$ and estimating the density function $f(\cdot)$ via kernel smoothing with a second-order kernel function at the optimal bandwidth rate, we have $u_n = O(n^{-1/2})$ and $v_n = O\{(n^{-1}\log n)^{2/5}\}$ from Proposition 2.1 and Hansen (2008), where $\{u_n, v_n\}$ are as defined in Assumption 2.2. These results on the convergence rates of $\{a_n, d_{n,2}, d_{n,\infty}, u_n, v_n\}$ are actually sufficient for the condition (2.11) in Theorem 2.1 and thus ensure the asymptotic normality of $\widehat{\theta}_{\text{SS}}$.

Remark 3.3 (Other reasonable choices of the nuisance estimator). As we conclude this section, we would like to reiterate that although we have focused on the combination of kernel smoothing and dimension reduction in this section, it is just one suitable strategy for approximating the nuisance function $\phi(\mathbf{X}, \theta)$ in (3.1). A wide class of alternatives could also be leveraged to estimate $\phi(\mathbf{X}, \theta)$ as long as the high-level conditions in Theorem 2.1 are satisfied. For instance, we can let \mathbf{P}_0 equal the $p \times p$ identity matrix in (3.1) and approximate $\phi(\mathbf{X}, \theta) \equiv \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$ by popular machine learning approaches, such as random forest (Breiman, 2001) and kernel machine regression (Liu et al., 2007), without use of dimension reduction. We will present the implementation details and numerical results related to random forest in Sections 4 and 5 while not delving any further into the theoretical aspects, which are beyond the main interest of this article.

4 Simulations

We study in this section the numerical performance of our proposed semi-supervised inference method on simulated data and compare it to the supervised counterpart. Throughout, the quantile level is $\tau = 0.5$. The results are similar for other quantile levels, such as $\tau = 0.25$ or 0.75 . We do not present them here for brevity. We set the sample sizes $n = 200, 500$ or $2,000$ and $N = 5,000$. The covariates \mathbf{X} are drawn from a p -dimensional normal distribution with a zero mean and an identity covariance matrix, where we choose $p = 10, 20, 200$ or 500 . The conditional outcome model is chosen as $Y \mid \mathbf{X} \sim \text{Normal}\{m(\mathbf{X}), 1\}$ for a variety of $m(\cdot)$ discussed below.

Let $\mathbf{X}_q := (\mathbf{X}_{[1]}, \dots, \mathbf{X}_{[q]})^T$, where $q = p$ when $p \in \{10, 20\}$, and $q = 5$ or $\lceil p^{1/2} \rceil$ when $p \in \{200, 500\}$, represents the (effective) sparsity (fully dense for $p \in \{10, 20\}$, and sparse or moderately dense for $p \in \{200, 500\}$, respectively) of the true conditional mean model $m(\mathbf{X})$, which we set as:

- (a) $m(\mathbf{X}) \equiv 0$, a null model;
- (b) $m(\mathbf{X}) \equiv \mathbf{1}_q^T \mathbf{X}_q$, a linear model;
- (c) $m(\mathbf{X}) \equiv \mathbf{1}_q^T \mathbf{X}_q + (\mathbf{1}_q^T \mathbf{X}_q)^2/q$, a single index model;
- (d) $m(\mathbf{X}) \equiv (\mathbf{1}_q^T \mathbf{X}_q) \{1 + 2(\mathbf{0}_{q-\lceil q/2 \rceil}^T, \mathbf{1}_{\lceil q/2 \rceil}^T) \mathbf{X}_q/q\}$, a double index model;
- (e) $m(\mathbf{X}) \equiv \mathbf{1}_q^T \mathbf{X}_q + \|\mathbf{X}_q\|^2/3$, a quadratic model.

These models generally represent a broad class of relation between Y and \mathbf{X} , containing commonly encountered linear and non-linear (quadratic and interaction) effects, in both low and high dimensional scenarios. In each configuration, estimation and inference results of our semi-supervised estimators are summarized from 500 replications. In the interest of space, the results for $p = 10$ or 200 are given in Appendix B.

For any kernel smoothing steps involved, we always use the second-order Gaussian kernel and select the bandwidths by maximizing the following cross-validated likelihood function:

$$L(b) := \prod_{i \in \mathcal{I}_k^-} \{\widehat{\phi}_{n,k}^{(-i,b)}(\mathbf{X}_i, \widehat{\theta}_{\text{INIT}}) + \tau\}^{I(Y_i < \widehat{\theta}_{\text{INIT}})} [1 - \{\widehat{\phi}_{n,k}^{(-i,b)}(\mathbf{X}_i, \widehat{\theta}_{\text{INIT}}) + \tau\}]^{I(Y_i \geq \widehat{\theta}_{\text{INIT}})},$$

where $\widehat{\phi}_{n,k}^{(-i,b)}(\cdot, \cdot)$ is a leave-one-out version of (3.2) constructed with the data $\{(Y_j, \mathbf{X}_j^T)^T : j \in \mathcal{I}_k^- \setminus \{i\}\}$ ($k = 1, \dots, \mathbb{K}$) and the bandwidth $h_n \equiv b$. All regularized approaches are based on the L_1 penalty with tuning parameters chosen by ten-fold cross validation. We set the number of folds in

Table 1: Simulation results of Section 4: Efficiencies of the semi-supervised estimators relative to the supervised estimator. The **boldface** in each case represents the best efficiency.

n	$m(\mathbf{X})$	$p = 20$					$p = 500, q = 5$				$p = 500, q = \lceil p^{1/2} \rceil$			
		KS ₁	KS ₂	PR	RF	ORE	KS ₁	KS ₂	PR	ORE	KS ₁	KS ₂	PR	ORE
200	(a)	0.90	0.93	0.86	0.94	1.00	0.95	0.68	0.99	1.00	0.95	0.68	0.99	1.00
	(b)	4.76	4.36	2.58	1.38	4.38	2.31	1.41	1.88	2.50	2.96	1.14	1.22	4.62
	(c)	4.23	4.03	2.74	1.37	4.36	1.78	1.02	1.44	2.21	1.95	1.07	1.24	4.60
	(d)	3.71	3.40	2.48	1.34	4.11	1.37	0.85	1.34	2.12	1.77	1.05	1.24	4.24
	(e)	2.45	2.31	1.99	1.38	4.74	1.90	1.24	1.56	2.65	1.20	0.91	1.17	4.97
500	(a)	0.98	0.95	0.98	0.98	1.00	0.96	0.91	0.99	1.00	0.96	0.91	0.99	1.00
	(b)	3.99	3.78	3.55	1.45	3.70	2.64	2.53	2.39	2.31	3.90	3.22	2.32	3.86
	(c)	3.76	3.96	3.54	1.43	3.69	2.13	1.94	1.68	2.07	3.56	3.05	2.35	3.84
	(d)	3.48	3.31	3.36	1.44	3.51	1.65	1.71	1.56	2.00	3.21	2.89	2.25	3.60
	(e)	2.38	2.32	2.24	1.47	3.93	1.70	1.63	1.68	2.43	2.00	1.89	1.58	4.09
2000	(a)	1.00	1.00	0.99	0.99	1.00	1.00	0.96	1.00	1.00	1.00	0.96	1.00	1.00
	(b)	2.65	2.66	2.61	1.45	2.34	1.76	1.75	1.74	1.80	2.29	2.28	2.04	2.39
	(c)	2.66	2.69	2.65	1.45	2.34	1.57	1.57	1.47	1.69	2.23	2.29	2.04	2.39
	(d)	2.59	2.62	2.58	1.46	2.28	1.48	1.56	1.42	1.65	2.15	2.16	1.98	2.31
	(e)	1.72	1.71	1.69	1.38	2.42	1.65	1.64	1.63	1.86	1.72	1.73	1.63	2.46

Glossary of notation: p , the dimension of \mathbf{X} ; q , the sparsity level; n , the labeled data size; $m(\mathbf{X}) \equiv \mathbb{E}(Y \mid \mathbf{X})$; KS₁/KS₂, kernel smoothing on the one/two direction(s) selected by linear regression/sliced inverse regression; RF, random forest; PR, parametric regression; ORE, oracle relative efficiency.

the cross-fitting process (2.6) as $\mathbb{K} = 10$. The initial estimator $\widehat{\theta}_{\text{INIT}}$ in (2.5) is chosen as the sample median while $\widehat{f}_n(\cdot)$ is taken as the kernel density estimator of Y , both obtained using \mathcal{L} .

To approximate the nuisance function $\phi(\mathbf{x}, \theta_0)$, the estimator $\widehat{\phi}_{n,k}(\mathbf{x}, \widehat{\theta}_{\text{INIT}})$ leveraging the data \mathcal{L}_k^- is calculated by:

- (I) *kernel smoothing* as in (3.2), where the $p \times r$ transformation matrix $\widehat{\mathbf{P}}_k$ is chosen as:
 - (i) the slope vector from unregularized or regularized linear regression ($r = 1$) of Y vs. \mathbf{X} ,
 - (ii) or the first two directions selected by the unregularized (with $\lceil n/5 \rceil$ slices of equal width) or regularized (with $\lceil n/75 \rceil$ slices of equal size) versions of sliced inverse regression ($r = 2$) (Li, 1991; Lin et al., 2019) of Y vs. \mathbf{X} ;
- (II) *parametric regression*, giving $\widehat{\phi}_{n,k}(\mathbf{x}, \widehat{\theta}_{\text{INIT}}) \equiv [1 + \exp\{-(\mathbf{1}, \mathbf{x}^T)\widehat{\gamma}_k\}]^{-1} - \tau$ with $\widehat{\gamma}_k$ being the slope vector from unregularized or regularized logistic regression of $I(Y < \widehat{\theta}_{\text{INIT}})$ vs. \mathbf{X} ;
- (III) *random forest* (for $p = 10$ or 20 only), treating $\psi(Y, \widehat{\theta}_{\text{INIT}})$ as the response, growing 500 trees and randomly sampling $\lceil p^{1/2} \rceil$ covariates as candidates at each split.

Here, regularization is applied when $p \in \{200, 500\}$. With these choices of $\widehat{\phi}_{n,k}(\mathbf{X}, \widehat{\theta}_{\text{INIT}})$, incorporating a variety of flexible and easy-to-implement (parametric, semi-parametric or nonparametric) approaches to fitting (working) models between a continuous or binary response and a set of possibly high dimensional covariates, we construct the nuisance estimators $\widehat{\phi}_n(\mathbf{X}, \widehat{\theta}_{\text{INIT}})$ via the cross-fitting process (2.6). Then our semi-supervised estimators $\widehat{\theta}_{\text{SS}}$ are obtained by plugging $\{\widehat{f}_n(\widehat{\theta}_{\text{INIT}}), \widehat{\phi}_n(\mathbf{X}, \widehat{\theta}_{\text{INIT}})\}$ in the one-step formula (2.5).

Table 2: Simulation results of Section 4: Inference based on the semi-supervised estimators using kernel smoothing on the direction selected by linear regression. All the numbers have been multiplied by 100. The **boldfaces** are the coverage rates of the 95% confidence intervals.

n	$m(\mathbf{X})$	$p = 20$				$p = 500, q = 5$				$p = 500, q = \lceil p^{1/2} \rceil$			
		ESE	Bias	ASE	CR	ESE	Bias	ASE	CR	ESE	Bias	ASE	CR
200	(a)	9.3	1.2	9.8	96.6	9.1	-0.7	9.6	96.6	9.1	-0.7	9.6	96.6
	(b)	19.0	-0.7	21.7	97.2	14.2	1.0	15.2	95.4	24.0	2.6	25.7	97.0
	(c)	19.9	-1.8	21.4	95.8	14.4	0.8	15.4	95.4	29.7	1.6	29.9	94.0
	(d)	20.2	4.1	22.0	97.2	15.6	0.7	18.8	98.2	28.1	1.2	30.5	96.6
	(e)	30.0	-4.3	32.6	95.2	17.4	-0.3	18.3	96.0	41.5	3.3	43.1	96.0
500	(a)	5.8	0.7	6.0	94.0	5.4	0.1	5.9	96.6	5.4	0.1	5.9	96.6
	(b)	12.8	-1.0	14.2	97.2	8.7	-0.1	9.6	97.0	14.0	2.1	15.1	95.8
	(c)	13.1	-1.5	13.8	95.8	8.3	0.7	9.0	96.4	14.6	1.7	15.4	95.2
	(d)	12.8	1.6	13.8	96.8	9.5	0.4	11.0	98.0	14.0	2.2	15.8	96.6
	(e)	18.5	-0.3	20.5	95.6	11.0	-0.3	11.2	95.6	21.7	1.5	22.6	97.0
2000	(a)	2.8	0.4	2.9	96.2	2.8	-0.1	2.9	95.0	2.8	-0.1	2.9	95.0
	(b)	8.4	-1.2	8.6	95.0	5.5	0.0	5.3	94.4	9.4	1.5	9.2	94.8
	(c)	8.2	-1.6	8.4	95.4	4.9	0.6	4.7	95.0	9.3	1.5	9.0	94.4
	(d)	7.8	0.6	8.1	96.2	5.1	0.0	5.5	95.2	8.6	2.0	8.9	95.8
	(e)	10.7	-1.8	11.0	95.6	5.9	-0.3	6.0	94.8	12.1	1.4	11.9	95.0

Glossary of notation: p , the dimension of \mathbf{X} ; q , the sparsity level; n , the labeled data size; $m(\mathbf{X}) \equiv \mathbb{E}(Y \mid \mathbf{X})$; ESE, empirical standard error; ASE, average of estimated standard errors; CR, coverage rate of 95% confidence intervals.

Table 1, along with Table 4 of Appendix B, presents the relative efficiencies, given by:

$$\mathbb{E}\{(\hat{\theta}_{\text{ss}} - \theta_0)^2\} / \mathbb{E}\{(\hat{\theta}_{\text{sup}} - \theta_0)^2\},$$

of our semi-supervised estimators to the supervised estimator, i.e., the sample median of the labeled data \mathcal{L} . For reference, we also provide the oracle relative efficiency $\sigma_{\text{sup}}^2 / \sigma_{\text{eff}}^2$ given by Proposition 2.1 and (2.13), which is achievable only asymptotically when the imputation function $\phi(\mathbf{X}, \theta) = \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{X}\}$. The true values of σ_{sup} , σ_{eff} and θ_0 are approximated by Monte Carlo based on 100,000 observations for $(Y, \mathbf{X}^T)^T$ independent of $\mathcal{L} \cup \mathcal{U}$. Except for the null model (a), where the unlabeled data \mathcal{U} does not help estimate θ_0 in theory, the various estimators based on kernel smoothing or random forest generally outperform the supervised method across all scenarios. These results coincide with the discussion on efficiency in Remark 2.5 considering that both kernel smoothing and random forest target the function $\mathbb{E}\{\psi(Y, \theta_0) \mid \mathbf{P}_0^T \mathbf{X}\}$ for some matrix \mathbf{P}_0 . Also, the parametric regression approach with imputation functions from logistic regression shows superiority over the sample median of \mathcal{L} under all the models other than (a), indicating that our approach is fully robust and that the logistic model captures a part of the relation between Y and \mathbf{X} . Moreover, we notice that the relative efficiencies of our estimators become closer to the corresponding oracle quantities as n increases, verifying the asymptotic optimality claimed in Remark 2.5. In summary, these observations demonstrate the efficiency gain achieved by our semi-supervised method relative to its supervised competitor.

Next, in Table 2 as well as Table 5 of Appendix B, we display, as a representative case, the results of inference concerning θ_0 based on our semi-supervised estimator $\hat{\theta}_{\text{ss}}$ with $\hat{\phi}_{n,k}(\cdot)$ constructed

using kernel smoothing on the direction selected by linear regression. We report the bias, the empirical standard error, the average of the estimated standard errors, and the coverage rate of the 95% confidence intervals, based on the asymptotic normality from Theorem 2.1. We can see that the biases are negligible, that the averages of the estimated standard errors are close to the corresponding empirical standard errors, and that the coverage rates are all around the nominal level of 0.95. Surprisingly, when the sample size $n = 200$, the dimension $p = 500$ and the sparsity level $q = \lceil p^{1/2} \rceil > n^{1/2}$, our method still generates satisfactory results and therefore shows insensitivity to the condition $\alpha_n = o(1)$ required by Theorem 3.1, considering $\alpha_n = q(\log p/n)^{1/2}$ in Assumption 3.1 when the L_1 penalty is leveraged under some suitable conditions (Bühlmann and Van De Geer, 2011; Negahban et al., 2012; Wainwright, 2019). Generally speaking, the numbers in these tables, which are yielded by the inference procedures based on the limiting distribution in Theorem 2.1 and the variance estimate in Remark 2.3, validate the theoretical results obtained in Section 2. With other choices of $\hat{\phi}_{n,k}(\cdot, \cdot)$, our method yields inference results similar in flavor to those in Tables 2 and 5. We thus omit them for the sake of brevity.

5 Real Data Analysis

In this section, we apply our semi-supervised method to a subset of data from the National Health and Nutrition Examination Survey Data I Epidemiologic Follow-up Study, a study jointly initiated by the National Center for Health Statistics and the National Institute on Aging in collaboration with other agencies of the United States Public Health Service (Hernán and Robins, 2020). The study aimed to explore the effect of various clinical, nutritional, demographic and behavioral factors on medical outcomes including morbidity and mortality. The data were collected through a baseline visit in 1971 and a follow-up visit in 1982, and the subset we focused on contained data on a cohort of 1425 individuals. A detailed description of the study data is available at <https://hspf.harvard.edu/miguel-hernan/causal-inference-book>.

Among the various biomedical outcomes recorded in the data, we are interested in the cohort's body weight at follow-up. Based on similar data, Ertefaie et al. (2022) considered the relationship between body weight and smoking from a causal perspective, estimating the average treatment effect of smoking cessation on weight gain during the follow up period. Unlike their study, which divided the observations into two groups according to smoking status (quit or not) at follow-up, we analyze all the 1,425 individuals together and our goal is to estimate the median weight θ_0 of the whole cohort in 1982. Further, we want to explore if there is any significant weight change among the cohort between 1971 and 1982, via comparing the analysis results to the baseline measure, i.e., the median weight 69.40, with a 95% confidence interval (68.32, 70.48), of these 1,425 individuals in 1971. Apart from body weight as the response, we also take into account $p = 20$ important covariates, to be included in our imputation models, whose names and descriptions are listed in Table 6 of Appendix C. To illustrate our approach in a semi-supervised setup, we randomly select $n = 100$ or 200 out of the 1,425 observations as the labeled data \mathcal{L} and regard the rest as the unlabeled data \mathcal{U} . Then we implement the supervised and semi-supervised strategies of estimation and inference as described in Section 4. Here, regularization is applied to all the regression procedures involved in our imputations. This process is replicated 500 times. We take the median weight $\hat{\theta}_{GS} = 72.12$, with an estimated standard error of 0.54, of all the 1,425 individuals in 1982 as a *gold standard* estimator. The subscript "GS" in $\hat{\theta}_{GS}$ stands for "gold standard". Summarized from the 500 replications, Table 3 reports the averages of the point estimates and 95%

Table 3: Data analysis results of Section 5: Estimation and inference of the cohort’s median weight in 1982 based on various methods, and efficiencies of the semi-supervised estimators relative to the supervised estimators, i.e., $\mathbb{E}\{(\hat{\theta}_{\text{SUP}} - \hat{\theta}_{\text{GS}})^2\}/\mathbb{E}\{(\hat{\theta}_{\text{SS}} - \hat{\theta}_{\text{GS}})^2\}$. The **boldface** in each case represents the best efficiency and the shortest confidence interval.

	$n = 100$					$n = 200$						
	Est	95% CI	RE	ESE	Bias	CR	Est	95% CI	RE	ESE	Bias	
Sup	71.85	(69.77, 73.93)	1.00	1.93	-0.27	0.96	71.78	(70.34, 73.22)	1.00	1.41	-0.34	0.95
SS-RF	72.14	(70.61, 73.68)	1.88	1.42	0.02	0.97	72.07	(71.02, 73.13)	2.06	1.01	-0.05	0.96
SS-KS ₁	72.15	(70.74, 73.56)	2.16	1.33	0.03	0.97	72.11	(71.12, 73.11)	2.57	0.90	-0.01	0.96
SS-KS ₂	72.17	(70.61, 73.73)	1.71	1.49	0.05	0.96	72.10	(71.07, 73.13)	2.33	0.95	-0.02	0.96
SS-PR	72.14	(70.69, 73.59)	2.01	1.37	0.02	0.97	72.11	(71.10, 73.12)	2.39	0.94	-0.01	0.96

Glossary of notation: n , the labeled data size; Est, point estimate; CI, confidence interval; RE, relative efficiency; ESE, empirical standard error; CR, coverage rate of the 95% confidence intervals; Sup, supervised estimator; SS, semi-supervised estimator; RF, random forest; KS₁/KS₂, kernel smoothing on the one/two direction(s) selected by linear regression/sliced inverse regression; PR, parametric regression.

confidence intervals, and the relative efficiencies $\mathbb{E}\{(\hat{\theta}_{\text{SUP}} - \hat{\theta}_{\text{GS}})^2\}/\mathbb{E}\{(\hat{\theta}_{\text{SS}} - \hat{\theta}_{\text{GS}})^2\}$ of the semi-supervised estimators versus the supervised estimators. We also present the inference results on θ_0 , where the biases and coverage rates are calculated relative to the gold standard, $\hat{\theta}_{\text{GS}}$.

In Table 3, which displays the analysis results on the cohort’s median weight in 1982, we notice that the lower bounds of all the semi-supervised confidence intervals are clearly above the upper bound of the 95% confidence interval (68.32, 70.48) of the median weight in 1971, indicating significant weight gain among the cohort between 1971 and 1982. However, this finding is likely to be ignored by the supervised method because the two supervised confidence intervals in Table 3 both overlap with (68.32, 70.48). This contrast demonstrates the considerable advantage of our semi-supervised inference procedures in terms of being more powerful in detecting significance. This is also reflected in their efficiencies, with the various semi-supervised estimators all yielding substantially better efficiencies (with relative efficiencies as high as 2.16 and 2.57, for the two choices of n) outperforming the supervised method. Moreover, they all have negligible biases (much lower than the supervised method) and also generate satisfactory coverage rates around the nominal level 0.95. These results confirm again the superiority of our semi-supervised method.

6 Discussion

We considered semi-supervised inference for quantile estimation in high dimensional settings, a problem relatively unaddressed in the existing literature, and developed a robust and efficient strategy based on imputation, allowing for flexible choices of the (nuisance) imputation model. We provided a complete characterization of the achievable estimators, and their robustness and efficiency properties. Moreover, we considered kernel smoothing estimators, with possible use of dimension reduction, as an illustration of the nuisance estimators involved in our method, establishing novel results on their uniform convergence rates in high dimensions. In a recent parallel work (Chakrabortty and Dai, 2022), we have also extended the theory and methodology developed in this article to the case of causal parameters. However, the problem of marginal quantiles itself without bringing in the causal framework is of general interest to the broader statistical community, so we focus on it in the current paper.

A natural extension of the marginal quantile estimation problem is quantile regression (Koenker, 2005) targeting the $(p+1)$ -dimensional parameter β_0 in a possibly misspecified *working* model that assumes the τ -conditional quantile of Y given \mathbf{X} equals $(1, \mathbf{X}^T)\beta_0$. One may expect that the one-step update strategy in Section 2.2 with suitable modifications can provide a family of semi-supervised estimators for β_0 that outperform the supervised counterpart. Such a procedure, however, also involves technical difficulties such as estimating the conditional density of $Y - (1, \mathbf{X}^T)\beta_0$ given \mathbf{X} , which can be a challenging task when p is moderate or large, and therefore requires developing more careful and sophisticated methodology. We thus leave this topic for future study.

Acknowledgment

AC's research was supported in part by the National Science Foundation grant NSF DMS-2113768. RJC's research was supported in part by a National Cancer Institute grant and by a Tripods grant from the National Science Foundation. RJC is also affiliated with the School of Mathematical and Physical Sciences, University of Technology Sydney, Broadway NSW 2007, Australia.

Supplementary material

- Appendices A–C include technical assumptions required by Theorem 3.1 (Appendix A.1), auxiliary lemmas that would be used for proving the main theorems (Appendix A.2), proofs of all the theoretical results (Appendices A.3–A.9), additional simulation results (Appendix B) and further information regarding the data analysis in Section 5 (Appendix C).
- All the computer programs used for obtaining the results in Sections 4–5 are available at: https://github.com/guorongdai/semi-supervised_quantile_estimation.

A Technical details

A.1 Smoothness conditions for kernel smoothing

The following Assumption A.1 contains the smoothness conditions required by Theorem 3.1. These conditions are fairly standard for kernel-based approaches, and their analogous versions can also be found in various existing works in the literature, such as Newey and McFadden (1994), Andrews (1995), Masry (1996) and Hansen (2008), among others.

Assumption A.1. (i) The function $K(\cdot)$ is a symmetric kernel of order $d \geq 2$ with a finite d th moment. Moreover, it is bounded, integrable and continuously differentiable. In addition, there exists some constant $v > 1$ such that $\|\nabla K(\mathbf{s})\| \leq c_1 \|\mathbf{s}\|^{-v}$ for any $\|\mathbf{s}\| > c_2$. (ii) The support \mathcal{S} of $\mathbf{S} \equiv \mathbf{P}_0^T \mathbf{X}$ is compact. The density function $f_{\mathbf{S}}(\cdot)$ of \mathbf{S} is bounded and bounded away from zero on \mathcal{S} . Further, it is d times continuously differentiable with a bounded d th derivative on \mathcal{S}_0 . (iii) With respect to \mathbf{s} , the conditional distribution function $F(\theta | \mathbf{S} = \mathbf{s})$ of Y given $\mathbf{S} = \mathbf{s}$ is d times continuously differentiable and has a bounded d th derivatives on $\mathcal{S}_0 \times \mathcal{B}(\theta_0, \varepsilon)$.

A.2 Preliminary lemmas

The following Lemmas A.1–A.2 would be useful in the proofs of the main theorems and propositions. The proofs of these lemmas, as well as Theorems 2.1–3.1 and Propositions 3.1–3.2, can be found in Sections A.3–A.9.

Lemma A.1. *For some fixed integer M , suppose $W_{n,1}, \dots, W_{n,M} \in \mathbb{R}$ are mutually independent sequences of random variables satisfying that, for some constants μ_m and $\sigma_m > 0$,*

$$W_{n,m} \xrightarrow{d} \text{Normal}(\mu_m, \sigma_m^2) \quad (m = 1, \dots, M; n \rightarrow \infty). \quad (\text{A.1})$$

Then, $\sum_{m=1}^M W_{n,m} \xrightarrow{d} \text{Normal}(\sum_{m=1}^M \mu_m, \sum_{m=1}^M \sigma_m^2)$ ($n \rightarrow \infty$).

Lemma A.2. *Suppose there are two independent samples, \mathcal{S}_1 and \mathcal{S}_2 , consisting of n and m independent copies of $(\mathbf{X}^T, Y)^T$, respectively. For $\gamma \in \mathbb{R}^d$ with some fixed d , let $\hat{g}_n(\mathbf{x}, \gamma)$ be an estimator of a measurable function $g(\mathbf{x}, \gamma) \in \mathbb{R}$ based on \mathcal{S}_1 and*

$$\mathbb{G}_m\{\hat{g}_n(\mathbf{X}, \gamma)\} := m^{1/2}[m^{-1} \sum_{(\mathbf{X}_i^T, Y_i)^T \in \mathcal{S}_2} \hat{g}_n(\mathbf{X}_i, \gamma) - \mathbb{E}_{\mathbf{X}}\{\hat{g}_n(\mathbf{X}, \gamma)\}].$$

For some set $\mathcal{T} \subset \mathbb{R}^d$, denote

$$\Delta(\mathcal{S}_1) := (\sup_{\gamma \in \mathcal{T}} \mathbb{E}_{\mathbf{X}}[\{\hat{g}_n(\mathbf{X}, \gamma)\}^2])^{1/2} \text{ and } M(\mathcal{S}_1) := \sup_{\mathbf{x} \in \mathcal{X}, \gamma \in \mathcal{T}} |\hat{g}_n(\mathbf{x}, \gamma)|.$$

For any $\eta \in (0, \Delta(\mathcal{S}_1) + c]$, suppose $\mathcal{G}_n := \{\hat{g}_n(\mathbf{X}, \gamma) : \gamma \in \mathcal{T}\}$ satisfies that

$$N_{[]} \{\eta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}})\} \leq H(\mathcal{S}_1) \eta^{-c}, \quad (\text{A.2})$$

with some function $H(\mathcal{S}_1) > 0$. Here \mathcal{G}_n is indexed by γ only and treats $\hat{g}_n(\cdot, \gamma)$ as a nonrandom function. Assume $H(\mathcal{S}_1) = O_p(a_n)$, $\Delta(\mathcal{S}_1) = O_p(d_{n,2})$ and $M(\mathcal{S}_1) = O_p(d_{n,\infty})$ with some positive sequences a_n , $d_{n,2}$ and $d_{n,\infty}$ allowed to diverge. Then, we have:

$$\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m\{\hat{g}_n(\mathbf{X}, \gamma)\}| = O_p(r_{n,m}),$$

where $r_{n,m} = d_{n,2}\{\log a_n + \log(d_{n,2}^{-1})\} + m^{-1/2}d_{n,\infty}\{(\log a_n)^2 + (\log d_{n,2})^2\}$.

A.3 Proof of Lemma A.1

We have that, for any $t \in \mathbb{R}$,

$$\begin{aligned} \mathbb{E}\{\exp(it \sum_{m=1}^M W_{n,m})\} &= \prod_{m=1}^M \mathbb{E}\{\exp(it W_{n,m})\} \\ &\rightarrow \prod_{m=1}^M \exp(i\mu_m t - \sigma_m^2 t^2/2) \\ &= \exp\{i(\sum_{m=1}^M \mu_m)t - (\sum_{m=1}^M \sigma_m^2)t^2/2\} \end{aligned} \quad (\text{A.3})$$

where i is the imaginary unit. In the above, the first step uses the mutual independence of $W_{n,1}, \dots, W_{n,M}$, and the second step is due to (A.1) and Levy's continuity theorem. The fact that (A.3) is the characteristic function of $\text{Normal}(\sum_{m=1}^M \mu_m, \sum_{m=1}^M \sigma_m^2)$ implies the conclusion. \square

A.4 Proof of Lemma A.2

For any $\delta \in (0, \Delta(\mathcal{S}_1) + c]$, we have that the bracketing integral:

$$\begin{aligned} J_{[]} \{ \delta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}}) \} &\equiv \int_0^\delta [1 + \log N_{[]} \{ \eta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}}) \}]^{1/2} d\eta \\ &\leq \int_0^\delta 1 + \log N_{[]} \{ \eta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}}) \} d\eta \\ &\leq \int_0^\delta 1 + \log H(\mathcal{S}_1) - c \log \eta d\eta \\ &= \delta \{ 1 + \log H(\mathcal{S}_1) \} + c(\delta - \delta \log \delta), \end{aligned}$$

where the third step is due to (A.2). This, combined with Lemma 19.36 of [Van der Vaart \(2000\)](#), implies that

$$\begin{aligned} &\mathbb{E}_{\mathbf{X}} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}|] \\ &\leq J_{[]} \{ \delta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}}) \} + [J_{[]} \{ \delta, \mathcal{G}_n \mid \mathcal{S}_1, L_2(\mathbb{P}_{\mathbf{X}}) \}]^2 M(\mathcal{S}_1) \delta^{-2} m^{-1/2} \\ &\leq \delta \{ 1 + \log H(\mathcal{S}_1) \} + c(\delta - \delta \log \delta) + \{ 1 + \log H(\mathcal{S}_1) + c(1 - \log \delta) \}^2 M(\mathcal{S}_1) m^{-1/2} \end{aligned}$$

for any $\delta \in (\Delta(\mathcal{S}_1), \Delta(\mathcal{S}_1) + c]$. Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}|] &\leq \Delta(\mathcal{S}_1) \{ 1 + \log H(\mathcal{S}_1) \} + c \{ \Delta(\mathcal{S}_1) - \Delta(\mathcal{S}_1) \log \Delta(\mathcal{S}_1) \} + \\ &\quad [1 + \log H(\mathcal{S}_1) + c \{ 1 - \log \Delta(\mathcal{S}_1) \}]^2 M(\mathcal{S}_1) m^{-1/2}. \end{aligned}$$

Since the right hand side in the above is $O_p(r_{n,m})$, it gives that

$$\mathbb{E}_{\mathbf{X}} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}|] = O_p(r_{n,m}). \quad (\text{A.4})$$

Then, for any positive sequence $t_n \rightarrow \infty$, we have:

$$\begin{aligned} &\mathbb{P} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}| > t_n r_{n,m} \mid \mathcal{S}_1] \\ &\leq (t_n r_{n,m})^{-1} \mathbb{E}_{\mathbf{X}} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}|] = o_p(1), \end{aligned}$$

where the first step holds by Markov's inequality and the last step is due to (A.4). This, combined with Lemma 6.1 of [Chernozhukov et al. \(2018\)](#), gives that

$$\mathbb{P} [\sup_{\gamma \in \mathcal{T}} |\mathbb{G}_m \{ \widehat{g}_n(\mathbf{X}, \gamma) \}| > t_n r_{n,m}] \rightarrow 0,$$

which completes the proof. \square

A.5 Proof of Theorem 2.1

Write

$$\widehat{\theta}_{\text{ss}} - \theta_0 = \{ S_1(\widehat{\theta}_{\text{INIT}}) - \theta_0 \} + \{ \widehat{f}_n(\widehat{\theta}_{\text{INIT}}) \}^{-1} \{ S_2(\widehat{\theta}_{\text{INIT}}) + S_3(\widehat{\theta}_{\text{INIT}}) \}, \quad (\text{A.5})$$

where

$$\begin{aligned} S_1(\theta) &:= \theta - \{ \widehat{f}_n(\theta) \}^{-1} \mathbb{E}_n \{ \psi(Y, \theta) \}, \\ S_2(\theta) &:= (1 - \nu_{n,N}) [\mathbb{E}_n \{ \widehat{\phi}_n(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta) \} - \mathbb{E}_N \{ \widehat{\phi}_n(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta) \}], \\ S_3(\theta) &:= \mathbb{E}_n \{ \phi(\mathbf{X}, \theta) \} - \mathbb{E}_{n+N} \{ \phi(\mathbf{X}, \theta) \}. \end{aligned}$$

Assumption 2.2 gives:

$$\mathbb{P}\{\widehat{\theta}_{\text{INIT}} \in \mathcal{B}(\theta_0, \varepsilon)\} \rightarrow 1, \quad (\text{A.6})$$

$$\widehat{L}_n := \{\widehat{f}_n(\widehat{\theta}_{\text{INIT}})\}^{-1} - \{f(\theta_0)\}^{-1} = O_p(v_n) = o_p(1). \quad (\text{A.7})$$

By Theorem 19.3 of [Van der Vaart \(2000\)](#), we know that $\{I(Y < \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}$ forms a \mathbb{P} -Donsker class, so the permanence properties of \mathbb{P} -Donsker classes ([Van der Vaart and Wellner, 1996](#)) guarantee that

$$\mathcal{D} := \{\psi(Y, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\} = \{I(Y < \theta) - \tau : \theta \in \mathcal{B}(\theta_0, \varepsilon)\} \quad (\text{A.8})$$

is also a \mathbb{P} -Donsker class. Moreover, the convergence (A.6) implies that $\psi(Y, \widehat{\theta}_{\text{INIT}})$ is in \mathcal{D} with probability tending to one. In addition, we have:

$$\begin{aligned} \mathbb{E}[\{\psi(Y, \widehat{\theta}_{\text{INIT}}) - \psi(Y, \theta_0)\}^2] &= \mathbb{E}[\{I(Y < \widehat{\theta}_{\text{INIT}}) - I(Y < \theta_0)\}^2] \\ &= F(\widehat{\theta}_{\text{INIT}}) + F(\theta_0) - 2F\{\min(\widehat{\theta}_{\text{INIT}}, \theta_0)\} \rightarrow 0. \end{aligned}$$

in probability because of the continuity of $F(\cdot)$ from Assumption 2.1 and the consistency of $\widehat{\theta}_{\text{INIT}}$ from Assumption 2.2. Hence Lemma 19.24 of [Van der Vaart \(2000\)](#) gives that

$$\mathbb{G}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}}) - \psi(Y, \theta_0)\} = o_p(1), \quad (\text{A.9})$$

which implies that

$$\begin{aligned} \mathbb{E}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} - \mathbb{E}\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} &= n^{-1/2}[\mathbb{G}_n\{\psi(Y, \theta_0)\} + \mathbb{G}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}}) - \psi(Y, \theta_0)\}] \\ &= \mathbb{E}_n\{\psi(Y, \theta_0)\} + o_p(n^{-1/2}). \end{aligned} \quad (\text{A.10})$$

Taylor's expansion gives that

$$\mathbb{E}\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} = f(\theta_0)(\widehat{\theta}_{\text{INIT}} - \theta_0) + O_p(|\widehat{\theta}_{\text{INIT}} - \theta_0|^2) \quad (\text{A.11})$$

$$= f(\theta_0)(\widehat{\theta}_{\text{INIT}} - \theta_0) + O_p(u_n^2) \quad (\text{A.12})$$

$$= O_p(u_n),$$

where the residual term in the first step is due to (A.6) as well as the fact that $f(\cdot)$ has a bounded derivative in $\mathcal{B}(\theta_0, \varepsilon)$ from Assumption 2.1, the second step uses Assumption 2.2, and the last step holds by the fact that $u_n = o(1)$ from Assumption 2.2. Then, we have:

$$\begin{aligned} \widehat{L}_n \mathbb{E}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} &= \widehat{L}_n[\mathbb{E}_n\{\psi(Y, \theta_0)\} + O_p(u_n) + o_p(n^{-1/2})] \\ &= \widehat{L}_n\{O_p(n^{-1/2}) + O_p(u_n) + o_p(n^{-1/2})\} \\ &= O_p(u_n v_n) + o_p(n^{-1/2}), \end{aligned} \quad (\text{A.13})$$

where the first step holds by (A.10) and (A.12), the second step uses the central limit theorem and the last step is due to (A.7). Thus, we have:

$$\begin{aligned} &S_1(\widehat{\theta}_{\text{INIT}}) - \theta_0 \\ &= \widehat{\theta}_{\text{INIT}} - \theta_0 - \{\widehat{f}_n(\widehat{\theta}_{\text{INIT}})\}^{-1} \mathbb{E}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} \\ &= \widehat{\theta}_{\text{INIT}} - \theta_0 - \{f(\theta_0)\}^{-1} \mathbb{E}_n\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} + O_p(u_n v_n) + o_p(n^{-1/2}) \\ &= \widehat{\theta}_{\text{INIT}} - \theta_0 - \{f(\theta_0)\}^{-1} [\mathbb{E}\{\psi(Y, \widehat{\theta}_{\text{INIT}})\} + \mathbb{E}_n\{\psi(Y, \theta_0)\}] + O_p(u_n v_n) + o_p(n^{-1/2}) \\ &= -\{f(\theta_0)\}^{-1} \mathbb{E}_n\{\psi(Y, \theta_0)\} + O_p(u_n^2 + u_n v_n) + o_p(n^{-1/2}), \end{aligned} \quad (\text{A.14})$$

where the second step uses (A.13), the third step holds by (A.10) and the last step is due to (A.11).

Moreover, denote

$$\mathbb{G}_{n_{\mathbb{K}},k}\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\} = n_{\mathbb{K}}^{1/2}[n_{\mathbb{K}}^{-1}\sum_{i \in \mathcal{I}_k} \widehat{\psi}_{n,k}(\mathbf{X}_i, \theta) - \mathbb{E}_{\mathbf{X}}\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}] \quad (k = 1, \dots, \mathbb{K}).$$

Considering Assumption 2.4, Lemma A.2 gives that

$$\begin{aligned} & \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{G}_{n_{\mathbb{K}},k}\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}| = O_p(r_n), \\ & \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{G}_N\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}| \\ & = O_p[d_{n,2}\{\log a_n + \log(d_{n,2}^{-1})\} + N^{-1/2}d_{n,\infty}\{(\log a_n)^2 + (\log d_{n,2})^2\}] \\ & = O_p(r_n) \quad (k = 1, \dots, \mathbb{K}). \end{aligned} \quad (\text{A.15})$$

Hence, using (A.6), we have that, with probability tending to one,

$$\begin{aligned} |S_2(\widehat{\theta}_{\text{INIT}})| & \leq \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{E}_n\{\widehat{\phi}_n(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta)\} - \mathbb{E}_N\{\widehat{\phi}_n(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta)\}| \\ & = \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{K}^{-1}\sum_{k=1}^{\mathbb{K}} [n_{\mathbb{K}}^{-1/2}\mathbb{G}_{n_{\mathbb{K}},k}\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\} - \\ & \quad N^{-1/2}\mathbb{G}_N\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}]| \\ & \leq \mathbb{K}^{-1}\sum_{k=1}^{\mathbb{K}} [n_{\mathbb{K}}^{-1/2}\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{G}_{n_{\mathbb{K}},k}\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}| + \\ & \quad N^{-1/2}\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{G}_N\{\widehat{\psi}_{n,k}(\mathbf{X}, \theta)\}|] = O_p(n^{-1/2}r_n). \end{aligned} \quad (\text{A.16})$$

In addition, we know

$$\widehat{f}_n(\widehat{\theta}_{\text{INIT}}) = O_p(1) \quad (\text{A.17})$$

due to the facts that $\widehat{f}_n(\widehat{\theta}_{\text{INIT}}) - f(\theta_0) = o_p(1)$ from Assumption 2.2, and that $f(\theta_0) > 0$ from Assumption 2.1. Combining (A.16) and (A.17) yields:

$$\{\widehat{f}_n(\widehat{\theta}_{\text{INIT}})\}^{-1}S_2(\widehat{\theta}_{\text{INIT}}) = O_p(n^{-1/2}r_n). \quad (\text{A.18})$$

Next, we have that

$$\begin{aligned} S_3(\widehat{\theta}_{\text{INIT}}) & = (\mathbb{E}_n - \mathbb{E}_{n+N})\{\phi(\mathbf{X}, \widehat{\theta}_{\text{INIT}})\} \\ & = (\mathbb{E}_n - \mathbb{E}_{n+N})\{\phi(\mathbf{X}, \theta_0)\} + n^{-1/2}\mathbb{G}_n\{\phi(\mathbf{X}, \widehat{\theta}_{\text{INIT}}) - \phi(\mathbf{X}, \theta_0)\} - \\ & \quad (n+N)^{-1/2}\mathbb{G}_{n+N}\{\phi(\mathbf{X}, \widehat{\theta}_{\text{INIT}}) - \phi(\mathbf{X}, \theta_0)\} \\ & = (\mathbb{E}_n - \mathbb{E}_{n+N})\{\phi(\mathbf{X}, \theta_0)\} + o_p(n^{-1/2}) + o_p\{(n+N)^{-1/2}\} \\ & = (\mathbb{E}_n - \mathbb{E}_{n+N})\{\phi(\mathbf{X}, \theta_0)\} + o_p(n^{-1/2}) \end{aligned} \quad (\text{A.19})$$

where the third step uses (A.6) and (2.7) in Assumption 2.3. Therefore, it follows that

$$\begin{aligned} \widehat{L}_n S_3(\widehat{\theta}_{\text{INIT}}) & = \widehat{L}_n[n^{-1/2}\mathbb{G}_n\{\phi(\mathbf{X}, \theta_0)\} - (n+N)^{-1/2}\mathbb{G}_{n+N}\{\phi(\mathbf{X}, \theta_0)\} + o_p(n^{-1/2})] \\ & = o_p(n^{-1/2}), \end{aligned} \quad (\text{A.20})$$

where the last step holds by (A.7) as well as the fact that $\mathbb{G}_n\{\phi(\mathbf{X}, \theta_0)\} = O_p(1)$ and $\mathbb{G}_{n+N}\{\phi(\mathbf{X}, \theta_0)\} = O_p(1)$ ensured by the central limit theorem and the square integrability of $\phi(\mathbf{X}, \theta_0)$ from Assumption 2.3. Combining (A.19) and (A.20) yields:

$$\begin{aligned} \{\widehat{f}_n(\widehat{\theta}_{\text{INIT}})\}^{-1}S_3(\widehat{\theta}_{\text{INIT}}) & = \{f(\theta_0)\}^{-1}S_3(\widehat{\theta}_{\text{INIT}}) + o_p(n^{-1/2}) \\ & = \{f(\theta_0)\}^{-1}(\mathbb{E}_n - \mathbb{E}_{n+N})\{\phi(\mathbf{X}, \theta_0)\} + o_p(n^{-1/2}). \end{aligned} \quad (\text{A.21})$$

Summing up, the equations (A.5), (A.14), (A.18) and (A.21) imply that

$$\hat{\theta}_{\text{ss}} - \theta_0 = \{f(\theta_0)\}^{-1} \mathbb{E}_n\{\omega_{n,N}(\mathbf{Z}, \theta_0)\} + O_p(u_n^2 + u_n v_n + n^{-1/2} r_n) + o_p(n^{-1/2}). \quad (\text{A.22})$$

Further, we know that

$$\begin{aligned} n^{1/2} \mathbb{E}_n\{\omega_{n,N}(\mathbf{Z}, \theta_0)\} &= \mathbb{G}_n\{(1 - \nu_{n,N})\phi(\mathbf{X}, \theta_0) - \psi(Y, \theta_0)\} - \\ &\quad (n/N)^{1/2} \mathbb{G}_N\{(1 - \nu_{n,N})\phi(\mathbf{X}, \theta_0)\}. \end{aligned} \quad (\text{A.23})$$

The central limit theorem and Slutsky's theorem give that, as $n, N \rightarrow \infty$,

$$\sigma_1^{-1} \mathbb{G}_n\{(1 - \nu_{n,N})\phi(\mathbf{X}, \theta_0) - \psi(Y, \theta_0)\} \rightarrow \text{Normal}(0, 1), \quad (\text{A.24})$$

$$\sigma_2^{-1} (n/N)^{1/2} \mathbb{G}_N\{(1 - \nu_{n,N})\phi(\mathbf{X}, \theta_0)\} \rightarrow \text{Normal}(0, 1), \quad (\text{A.25})$$

where

$$\begin{aligned} \sigma_1^2 &:= \mathbb{E}[\{\psi(Y, \theta_0)\}^2] + (1 - \nu_{n,N})^2 \text{var}\{\phi(\mathbf{X}, \theta_0)\} - 2(1 - \nu_{n,N}) \mathbb{E}\{\psi(Y, \theta_0)\phi(\mathbf{X}, \theta_0)\}, \\ \sigma_2^2 &:= (n/N)(1 - \nu_{n,N})^2 \text{var}\{\phi(\mathbf{X}, \theta_0)\}. \end{aligned}$$

Thus, we have:

$$\begin{aligned} \sigma_1^2 + \sigma_2^2 &= \mathbb{E}[\{\psi(Y, \theta_0)\}^2] + (1 - \nu_{n,N}) \text{var}\{\phi(\mathbf{X}, \theta_0)\} - 2(1 - \nu_{n,N}) \mathbb{E}\{\phi(\mathbf{X}, \theta_0)\psi(Y, \theta_0)\} \\ &= (1 - \nu_{n,N}) \text{var}\{\psi(Y, \theta_0) - \phi(\mathbf{X}, \theta_0)\} + \nu_{n,N} \text{var}\{\psi(Y, \theta_0)\} = \sigma_{\text{ss}}^2. \end{aligned} \quad (\text{A.26})$$

Finally, applying Lemma A.1 and Slutsky's theorem, the equations (A.22)–(A.26) conclude the asymptotic normality under the assumption (2.11) and the independence of the empirical processes in (A.24) and (A.25). \square

A.6 Proof of Corollary 2.1

Since $\nu = 0$, the central limit theorem gives that

$$\mathbb{E}_{n+N}\{\phi(\mathbf{X}, \theta_0)\} = \mathbb{E}\{\phi(\mathbf{X}, \theta_0)\} + O_p\{(n + N)^{-1/2}\} = \mathbb{E}\{\phi(\mathbf{X}, \theta_0)\} + o_p(n^{-1/2}).$$

This, combined with (A.22), implies the stochastic expansion, followed by the asymptotic normality. Further, it is clear that $\sigma_{\text{ss}} \rightarrow \tilde{\sigma}_{\text{ss}}$ as $n \rightarrow \infty$ in that $\lim_{n \rightarrow \infty} \nu_{n,N} = 0$. \square

A.7 Proof of Proposition 3.1

The second moment of $\phi(\mathbf{X}, \theta_0)$ is obviously finite because the function $F(\cdot \mid \mathbf{S})$ is bounded. For any $\theta_1, \theta_2 \in \mathcal{B}(\theta_0, \varepsilon)$, Taylor's expansion gives:

$$\begin{aligned} |\phi(\mathbf{X}, \theta_1) - \phi(\mathbf{X}, \theta_2)| &= |F(\theta_1 \mid \mathbf{S}) - F(\theta_2 \mid \mathbf{S})| \\ &\leq \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} f(\theta \mid \mathbf{S}) |\theta_1 - \theta_2|, \end{aligned}$$

Therefore, under the assumption that

$$\mathbb{E}[\{\sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} f(\theta \mid \mathbf{S})\}^2] < \infty, \quad (\text{A.27})$$

Example 19.7 of [Van der Vaart \(2000\)](#) implies:

$$N_{[]} \{ \eta, \mathcal{F}, L_2(\mathbb{P}_{\mathbf{X}}) \} \leq c \eta^{-1} \quad (\text{A.28})$$

with $\mathcal{F} := \{ \phi(\mathbf{X}, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon) \}$. Then, by Theorem 19.5 of [Van der Vaart \(2000\)](#), we know that \mathcal{F} is \mathbb{P} -Donsker. Further, we have that, for any sequence $\tilde{\theta} \rightarrow \theta_0$ in probability,

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} [\{ \phi(\mathbf{X}, \tilde{\theta}) - \phi(\mathbf{X}, \theta_0) \}^2] &= \mathbb{E}_{\mathbf{S}} [\{ F(\tilde{\theta} | \mathbf{S}) - F(\theta_0 | \mathbf{S}) \}^2] \\ &\leq (\tilde{\theta} - \theta_0)^2 \mathbb{E} [\{ \sup_{\theta \in \mathcal{B}(\theta_0, \varepsilon)} f(\theta | \mathbf{S}) \}^2] \rightarrow 0 \end{aligned}$$

in probability, where the second step uses Taylor's expansion and the fact that $\tilde{\theta} \in \mathcal{B}(\theta_0, \varepsilon)$ with probability approaching one, and the last step holds by (A.27). Lastly, applying Lemma 19.24 of [Van der Vaart \(2000\)](#) concludes (2.7). \square

A.8 Proof of Theorem 3.1

Set $m(\mathbf{x}, \theta, \mathbf{P}) := \phi(\mathbf{x}, \theta, \mathbf{P}) f_{\mathbf{S}}(\mathbf{P}^T \mathbf{x})$. We now derive the convergence rate of $\hat{m}_{n,k}(\mathbf{x}, \theta, \hat{\mathbf{P}}_k) - m(\mathbf{x}, \theta, \mathbf{P}_0)$.

We first handle the error from estimating \mathbf{P}_0 by $\hat{\mathbf{P}}_k$, i.e., $\hat{m}_{n,k}(\mathbf{x}, \theta, \hat{\mathbf{P}}_k) - \hat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0)$. Taylor's expansion gives that, for

$$\bar{\mathbf{s}}_n := h_n^{-1} \{ \mathbf{P}_0^T + \mathbf{M}(\hat{\mathbf{P}}_k - \mathbf{P}_0)^T \} (\mathbf{x} - \mathbf{X}) \quad (\text{A.29})$$

with some $\mathbf{M} := \text{diag}(\mu_1, \dots, \mu_r)$ and $\mu_j \in (0, 1)$ ($j = 1, \dots, r$),

$$\begin{aligned} &\hat{m}_{n,k}(\mathbf{x}, \theta, \hat{\mathbf{P}}_k) - \hat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0) \\ &= h_n^{-(r+1)} \mathbb{E}_{n,k} [\{ \nabla K(\bar{\mathbf{s}}_n) \}^T (\hat{\mathbf{P}}_k - \mathbf{P}_0)^T (\mathbf{x} - \mathbf{X}) \psi(Y, \theta)] \\ &= h_n^{-(r+1)} \text{trace} ((\hat{\mathbf{P}}_k - \mathbf{P}_0)^T \mathbb{E}_{n,k} [(\mathbf{x} - \mathbf{X}) \{ \nabla K(\bar{\mathbf{s}}_n) \}^T \psi(Y, \theta)]) \\ &= h_n^{-(r+1)} \text{trace} [(\hat{\mathbf{P}}_k - \mathbf{P}_0)^T \{ \mathbf{U}_{n,1}(\mathbf{x}, \theta) + \mathbf{U}_{n,2}(\mathbf{x}, \theta) - \mathbf{U}_{n,3}(\mathbf{x}, \theta) \}], \end{aligned} \quad (\text{A.30})$$

where

$$\begin{aligned} \mathbf{U}_{n,1}(\mathbf{x}, \theta) &:= \mathbb{E}_{n,k} ((\mathbf{x} - \mathbf{X}) [\nabla K(\bar{\mathbf{s}}_n) - \nabla K\{h_n^{-1} \mathbf{P}_0^T (\mathbf{x} - \mathbf{X})\}]^T \psi(Y, \theta)), \\ \mathbf{U}_{n,2}(\mathbf{x}, \theta) &:= \mathbb{E}_{n,k} (\mathbf{x} [\nabla K\{h_n^{-1} \mathbf{P}_0^T (\mathbf{x} - \mathbf{X})\}]^T \psi(Y, \theta)), \\ \mathbf{U}_{n,3}(\mathbf{x}, \theta) &:= \mathbb{E}_{n,k} (\mathbf{X} [\nabla K\{h_n^{-1} \mathbf{P}_0^T (\mathbf{x} - \mathbf{X})\}]^T \psi(Y, \theta)). \end{aligned}$$

For the function $\rho(\cdot)$ in Assumption 3.2 (ii), denote $\mathcal{J}_n := \{ h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\} : \mathbf{s} \in \mathcal{S} \}$. Taylor's expansion gives that, for any $\mathbf{s}_1, \mathbf{s}_2 \in \mathcal{S}$ and some $\bar{\mathbf{s}} := \mathbf{s}_1 + \mathbf{M}(\mathbf{s}_2 - \mathbf{s}_1)$ with $\mathbf{M} := \text{diag}(\mu_1, \dots, \mu_r)$ and $\mu_j \in (0, 1)$ ($j = 1, \dots, r$),

$$\begin{aligned} &h_n^{-r} |\rho\{h_n^{-1}(\mathbf{s}_1 - \mathbf{S})\} - \rho\{h_n^{-1}(\mathbf{s}_2 - \mathbf{S})\}| \\ &= h_n^{-(r+1)} |[\nabla \rho\{h_n^{-1}(\bar{\mathbf{s}} - \mathbf{S})\}]^T (\mathbf{s}_1 - \mathbf{s}_2)| \leq c h_n^{-(r+1)} \|\mathbf{s}_1 - \mathbf{s}_2\|, \end{aligned}$$

where the second step uses the boundedness of $\nabla \rho(\cdot)$ from Assumption 3.2 (ii). Therefore Example 19.7 of [Van der Vaart \(2000\)](#) implies:

$$N_{[]} \{ \eta, \mathcal{J}_n, L_2(\mathbb{P}_{\mathbf{X}}) \} \leq c h_n^{-(r+1)} \eta^{-r}. \quad (\text{A.31})$$

Moreover, we have that

$$\sup_{\mathbf{s}, \mathbf{S} \in \mathcal{S}} [h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}] = O(h_n^{-r}). \quad (\text{A.32})$$

due to the boundedness of $\rho(\cdot)$ from Assumption 3.2 (ii). In addition, we know that

$$\begin{aligned} \sup_{\mathbf{s} \in \mathcal{S}} \mathbb{E}_{\mathbf{S}}([h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}]^2) &= h_n^{-r} \sup_{\mathbf{s} \in \mathcal{S}} \int h_n^{-r} [\rho\{h_n^{-1}(\mathbf{s} - \mathbf{v})\}]^2 f_{\mathbf{S}}(\mathbf{v}) d\mathbf{v} \\ &= h_n^{-r} \sup_{\mathbf{s} \in \mathcal{S}} \int \{\rho(\mathbf{t})\}^2 f_{\mathbf{S}}(\mathbf{s} - h_n \mathbf{t}) d\mathbf{t} = O(h_n^{-r}), \end{aligned} \quad (\text{A.33})$$

where the second step uses change of variables and the last step holds by the boundedness of $f_{\mathbf{S}}(\cdot)$ from Assumption A.1 (ii) and the square integrability of $\rho(\cdot)$ from Assumption 3.2 (ii). Based on (A.31)–(A.33), applying Lemma A.2 yields that

$$\begin{aligned} &\sup_{\mathbf{s} \in \mathcal{S}} |\mathbb{E}_{n,k} [h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}] - \mathbb{E}_{\mathbf{X}} [h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}]| \\ &= O_p\{n_{\mathbb{K}^-}^{-1/2} h_n^{-r/2} \log(h_n^{-1}) + n_{\mathbb{K}^-}^{-1} h_n^{-r} (\log h_n)^2\} = o_p(1), \end{aligned} \quad (\text{A.34})$$

where the second step is because we assume $(nh_n^r)^{-1/2} \log(h_n^{-r}) = o(1)$. Then, we know

$$\begin{aligned} \sup_{\mathbf{s} \in \mathcal{S}} \mathbb{E}_{\mathbf{S}} [h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}] &= \sup_{\mathbf{s} \in \mathcal{S}} \int h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{v})\} f_{\mathbf{S}}(\mathbf{v}) d\mathbf{v} \\ &= \sup_{\mathbf{s} \in \mathcal{S}} \int \rho(\mathbf{t}) f_{\mathbf{S}}(\mathbf{s} - h_n \mathbf{t}) d\mathbf{t} = O(1). \end{aligned}$$

where the second step uses change of variables and the last step holds by the boundedness of $f_{\mathbf{S}}(\cdot)$ from Assumption A.1 (ii) and the integrability of $\rho(\cdot)$ from Assumption 3.2 (ii). This, combined with (A.34), implies that

$$\sup_{\mathbf{s} \in \mathcal{S}} \mathbb{E}_{n,k} [h_n^{-r} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}] = O_p(1). \quad (\text{A.35})$$

Next, we have:

$$\begin{aligned} &\sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{n,k} [\|\nabla K(\bar{\mathbf{s}}_n) - \nabla K\{h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\}\|] \\ &\leq \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{n,k} [\|\bar{\mathbf{s}}_n - h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\| \rho\{h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\}] \\ &\leq \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{n,k} [\|(\widehat{\mathbf{P}}_k - \mathbf{P}_0)^T(\mathbf{x} - \mathbf{X})\| h_n^{-1} \rho\{h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\}] \\ &\leq c \|\widehat{\mathbf{P}}_k - \mathbf{P}_0\|_1 \sup_{\mathbf{x}, \mathbf{X} \in \mathcal{X}} \|\mathbf{x} - \mathbf{X}\|_{\infty} \sup_{\mathbf{s} \in \mathcal{S}} \mathbb{E}_{n,k} [h_n^{-1} \rho\{h_n^{-1}(\mathbf{s} - \mathbf{S})\}] \\ &= O_p(h_n^{r-1} \alpha_n), \end{aligned} \quad (\text{A.36})$$

where the first step uses the local lipschitz continuity of $\nabla K(\cdot)$ from Assumption 3.2 (ii), the second step is due to the definition (A.29) of $\bar{\mathbf{s}}_n$, the third step holds by Hölder's inequality, and the last step is because of Assumptions 3.1, 3.2 (i) and the equation (A.35). Hence,

$$\begin{aligned} &\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,1}(\mathbf{x}, \theta)\|_{\infty} \\ &\leq c \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{n,k} [\|\mathbf{x} - \mathbf{X}\|_{\infty} \|\nabla K(\bar{\mathbf{s}}_n) - \nabla K\{h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\}\|] \\ &\leq c \sup_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{n,k} [\|\nabla K(\bar{\mathbf{s}}_n) - \nabla K\{h_n^{-1}\mathbf{P}_0^T(\mathbf{x} - \mathbf{X})\}\|] = O_p(h_n^{r-1} \alpha_n). \end{aligned}$$

where the first step holds by the boundedness of $\psi(Y, \theta)$, the second step is due to Assumption 3.2 (i), and the last step uses (A.36). This, combined with Assumption 3.1 and Hölder's inequality, implies that

$$\begin{aligned} &\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|(\widehat{\mathbf{P}}_k - \mathbf{P}_0)^T \mathbf{U}_{n,1}(\mathbf{x}, \theta)\|_{\infty} \\ &\leq \|\widehat{\mathbf{P}}_k - \mathbf{P}_0\|_1 \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,1}(\mathbf{x}, \theta)\|_{\infty} = O_p(h_n^{r-1} \alpha_n^2). \end{aligned} \quad (\text{A.37})$$

Next, under Assumptions A.1 (ii) and 3.2 (ii) as long as the fact that $\{\psi(Y, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}$ is a VC class with a bounded envelope function, Lemma B.4 of Escanciano et al. (2014) gives that

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,2}(\mathbf{x}, \theta) - \mathbb{E}\{\mathbf{U}_{n,2}(\mathbf{x}, \theta)\}\|_\infty = O_p(h_n^r \gamma_n), \quad (\text{A.38})$$

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,3}(\mathbf{x}, \theta) - \mathbb{E}\{\mathbf{U}_{n,3}(\mathbf{x}, \theta)\}\|_\infty = O_p(h_n^r \gamma_n). \quad (\text{A.39})$$

Let $\delta(\mathbf{s}, \theta) := \mathbb{E}\{\psi(Y, \theta) \mid \mathbf{S} = \mathbf{s}\} f_{\mathbf{S}}(\mathbf{s})$ and $\nabla \delta(\mathbf{s}, \theta) := \partial \delta(\mathbf{s}, \theta) / \partial \mathbf{s}$. We have:

$$\begin{aligned} & \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbb{E}\{\mathbf{U}_{n,2}(\mathbf{x}, \theta)\}\|_\infty \\ & \leq \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{x} \int \delta(\mathbf{s}, \theta) [\nabla K\{h_n^{-1}(\mathbf{P}_0^T \mathbf{x} - s)\}]^T ds\|_\infty \\ & = h^{r+1} \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{x} \int \{\nabla \delta(\mathbf{P}_0^T \mathbf{x} - h_n \mathbf{t}, \theta)\}^T K(\mathbf{t}) d\mathbf{t}\|_\infty = O(h^{r+1}). \end{aligned} \quad (\text{A.40})$$

In the above, the second step uses integration by parts and change of variables, while the last step holds by Assumption 3.2 (i), the boundedness of $\nabla \delta(\mathbf{s}, \theta)$ from Assumptions A.1 (ii) and (iii), as well as the integrability of $K(\cdot)$ from Assumption A.1 (i). Set $\zeta(\mathbf{s}, \theta) := f_{\mathbf{S}}(\mathbf{s}) \boldsymbol{\eta}_1(\mathbf{s}, \theta)$ and $\nabla \zeta(\mathbf{s}, \theta) := \partial \zeta(\mathbf{s}, \theta) / \partial \mathbf{s}$. Analogous to (A.40), we know

$$\begin{aligned} & \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbb{E}\{\mathbf{U}_{n,3}(\mathbf{x}, \theta)\}\|_\infty \\ & \leq \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\int \zeta(\mathbf{s}, \theta) [\nabla K\{h_n^{-1}(\mathbf{P}_0^T \mathbf{x} - s)\}]^T ds\|_\infty \\ & = h^{r+1} \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\int \{\nabla \zeta(\mathbf{P}_0^T \mathbf{x} - h_n \mathbf{t}, \theta)\}^T K(\mathbf{t}) d\mathbf{t}\|_\infty = O(h^{r+1}), \end{aligned} \quad (\text{A.41})$$

where the last step holds by the boundedness of $\|\nabla \zeta(\mathbf{s}, \theta)\|_\infty$ from Assumptions A.1 (ii) and 3.2 (iii), and the integrability of $K(\cdot)$ from Assumption A.1 (i). Combining (A.38)–(A.41) yields:

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,2}(\mathbf{x}, \theta) - \mathbf{U}_{n,3}(\mathbf{x}, \theta)\|_\infty = O_p(h_n^r \gamma_n + h_n^{r+1}),$$

which implies that

$$\begin{aligned} & \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|(\mathbf{P}_0 - \widehat{\mathbf{P}}_k)^T \{\mathbf{U}_{n,2}(\mathbf{x}, \theta) - \mathbf{U}_{n,3}(\mathbf{x}, \theta)\}\|_\infty \\ & \leq \|\mathbf{P}_0 - \widehat{\mathbf{P}}_k\|_1 \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} \|\mathbf{U}_{n,2}(\mathbf{x}, \theta) - \mathbf{U}_{n,3}(\mathbf{x}, \theta)\|_\infty \\ & = O_p(h_n^r \gamma_n \alpha_n + h_n^{r+1} \alpha_n) \end{aligned}$$

using Hölder's inequality and Assumption 3.1. This, combined with (A.30) and (A.37), gives:

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - \widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0)| = O_p(s_{n,2}). \quad (\text{A.42})$$

Moreover, we control the error $\widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0) - m(\mathbf{x}, \theta, \mathbf{P}_0)$. Under Assumptions A.1 (i), (ii) and the fact that $\{\psi(Y, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}$ is a VC class with a bounded envelope function, Lemma B.4 of Escanciano et al. (2014) gives that

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0) - \mathbb{E}\{\widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0)\}| = O_p(\gamma_n). \quad (\text{A.43})$$

Further, under Assumption A.1, standard arguments based on d th order Taylor's expansion of $m(\mathbf{x}, \theta, \mathbf{P}_0)$ yield that

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\mathbb{E}\{m_{n,k}(\mathbf{x}, \theta, \mathbf{P}_0)\} - m(\mathbf{x}, \theta, \mathbf{P}_0)| = O(h_n^d). \quad (\text{A.44})$$

Combining (A.42)–(A.44) yields:

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - m(\mathbf{x}, \theta, \mathbf{P}_0)| = O_p(s_{n,1} + s_{n,2}). \quad (\text{A.45})$$

Similar arguments give that

$$\sup_{\mathbf{x} \in \mathcal{X}} |\widehat{\ell}_{n,k}(\mathbf{x}, \widehat{\mathbf{P}}_k) - f_{\mathbf{S}}(\mathbf{P}_0^T \mathbf{x})| = O_p(s_{n,1} + s_{n,2}). \quad (\text{A.46})$$

Thus, we have:

$$\begin{aligned} & \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\widehat{\phi}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - \phi(\mathbf{x}, \theta, \mathbf{P}_0)| \\ &= \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\{\widehat{\ell}_{n,k}(\mathbf{x}, \widehat{\mathbf{P}}_k)\}^{-1} \widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - \{\ell(\mathbf{x}, \mathbf{P}_0)\}^{-1} m(\mathbf{x}, \theta, \mathbf{P}_0)| \\ &\leq \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |\{\widehat{\ell}_{n,k}(\mathbf{x}, \mathbf{P}_0)\}^{-1} \{\widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k) - m(\mathbf{x}, \theta, \mathbf{P}_0)\}| + \\ & \quad \sup_{\mathbf{x} \in \mathcal{X}, \theta \in \mathcal{B}(\theta_0, \varepsilon)} |[\{\widehat{\ell}_{n,k}(\mathbf{x}, \mathbf{P}_0)\}^{-1} - \{\ell(\mathbf{x}, \mathbf{P}_0)\}^{-1}] m(\mathbf{x}, \theta, \mathbf{P}_0)| = O_p(s_{n,1} + s_{n,2}), \end{aligned}$$

where the last step follows from the fact that $O_p(s_{n,1} + s_{n,2}) = o(1)$ and repeated use of (A.45), (A.46) as well as Assumption A.1 (ii). \square

A.9 Proof of Proposition 3.2

Considering that

$$\widehat{\phi}_{n,k}(\mathbf{X}, \theta, \widehat{\mathbf{P}}_k) \equiv \{\widehat{\ell}_{n,k}(\mathbf{x}, \widehat{\mathbf{P}}_k)\}^{-1} \widehat{m}_{n,k}(\mathbf{x}, \theta, \widehat{\mathbf{P}}_k)$$

with $\widehat{m}_{n,k}(\mathbf{x}, \theta, \mathbf{P}) \equiv h_n^{-r} \mathbb{E}_{n,k}[\{I(Y < \theta) - \tau\} K_h\{\mathbf{P}^T(\mathbf{x} - \mathbf{X})\}]$, it is obvious that, given \mathcal{L} ,

$$\{\widehat{\phi}_{n,k}(\mathbf{X}, \theta, \widehat{\mathbf{P}}_k) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\} \subset \{\widehat{\phi}_{n,k}(\mathbf{X}, \theta_i, \widehat{\mathbf{P}}_k) : i = 1, \dots, n+1\}$$

for any $\theta_1 < Y_{(1)}$, $\theta_i \in [Y_{(i-1)}, Y_{(i)})$ ($i = 2, \dots, n$) and $\theta_{n+1} \geq Y_{(n)}$, where $Y_{(i)}$ is the i th order statistic of $\{Y_i : i = 1, \dots, n\}$. Therefore the set $\{\widehat{\phi}_{n,k}(\mathbf{X}, \theta, \widehat{\mathbf{P}}_k) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}$ contains at most $(n+1)$ different functions indexed by θ given \mathcal{L} . This, combined with (A.28), implies that the set $\mathcal{P}_{n,k} \equiv \{\widehat{\phi}_{n,k}(\mathbf{X}, \theta) - \phi(\mathbf{X}, \theta) : \theta \in \mathcal{B}(\theta_0, \varepsilon)\}$ satisfies:

$$N_{[]} \{\eta, \mathcal{P}_{n,k} \mid \mathcal{L}, L_2(\mathbb{P}_{\mathbf{X}})\} \leq c(n+1)\eta^{-1}. \quad \square$$

B Additional simulation results

We present here the results of our simulation studies with $p = 10$ or 200 , in Tables 4 and 5. See Section 4 for the descriptions of the settings and the methods. The behavior of the results – both in estimation and inference – are similar to the other cases presented in the main paper.

C Supplement to the data analysis in Section 5

The following Table 6 lists the names and descriptions of the covariates we considered for our data analysis in Section 5. These were the covariates included in all our imputation models.

Table 4: Simulation results of Section 4: Efficiencies of the semi-supervised estimators relative to the supervised estimator. The **boldface** in each case represents the best efficiency.

n	$m(\mathbf{X})$	$p = 10$					$p = 200, q = 5$				$p = 200, q = \lceil p^{1/2} \rceil$			
		KS ₁	KS ₂	PR	RF	ORE	KS ₁	KS ₂	PR	ORE	KS ₁	KS ₂	PR	ORE
200	(a)	0.94	0.91	0.93	0.94	1.00	0.96	0.81	0.98	1.00	0.96	0.81	0.98	1.00
	(b)	3.42	3.25	3.07	1.63	3.33	2.54	2.29	2.18	2.52	3.24	2.22	1.81	3.90
	(c)	2.98	2.88	2.84	1.55	3.21	2.02	1.80	1.58	2.23	2.62	2.13	1.81	3.86
	(d)	2.51	2.66	2.34	1.54	2.88	1.45	1.25	1.42	2.14	2.20	1.77	1.65	3.41
	(e)	2.20	2.16	1.94	1.56	3.58	1.80	1.66	1.58	2.67	1.75	1.43	1.44	4.19
500	(a)	0.98	0.99	0.96	0.96	1.00	0.97	0.93	0.99	1.00	0.97	0.93	0.99	1.00
	(b)	3.19	3.12	3.10	1.75	2.95	1.98	1.90	1.92	2.33	3.19	3.00	2.61	3.37
	(c)	2.86	2.87	2.85	1.67	2.86	1.94	1.86	1.65	2.09	2.94	3.02	2.59	3.34
	(d)	2.52	2.51	2.37	1.63	2.62	1.39	1.41	1.38	2.01	2.40	2.26	2.04	3.01
	(e)	2.13	2.08	2.12	1.74	3.14	2.01	1.94	1.88	2.45	2.24	2.12	2.00	3.57
2000	(a)	0.98	1.01	0.99	0.97	1.00	0.99	1.00	1.00	1.00	0.99	1.00	1.00	1.00
	(b)	1.95	1.96	1.94	1.58	2.08	1.89	1.89	1.85	1.81	2.19	2.17	2.10	2.23
	(c)	1.93	1.92	1.91	1.54	2.05	1.68	1.67	1.51	1.70	2.11	2.11	2.03	2.22
	(d)	1.74	1.76	1.70	1.47	1.94	1.46	1.61	1.43	1.65	2.09	2.13	1.99	2.10
	(e)	1.67	1.68	1.66	1.69	2.15	1.71	1.72	1.69	1.87	1.61	1.60	1.57	2.30

Glossary of notation: p , the dimension of \mathbf{X} ; q , the sparsity level; n , the labeled data size; $m(\mathbf{X}) \equiv \mathbb{E}(Y \mid \mathbf{X})$; KS₁/KS₂, kernel smoothing on the one/two direction(s) selected by linear regression/sliced inverse regression; RF, random forest; PR, parametric regression; ORE, oracle relative efficiency.

Table 5: Simulation results of Section 4: Inference based on the semi-supervised estimators using kernel smoothing on the direction selected by linear regression. All the numbers have been multiplied by 100. The **boldfaces** are the coverage rates of the 95% confidence intervals.

n	$m(\mathbf{X})$	$p = 10$				$p = 200, q = 5$				$p = 200, q = \lceil p^{1/2} \rceil$			
		ESE	Bias	ASE	CR	ESE	Bias	ASE	CR	ESE	Bias	ASE	CR
200	(a)	9.0	0.3	9.7	97.2	8.8	-0.2	9.6	96.6	8.8	-0.2	9.6	96.6
	(b)	17.3	0.1	17.6	93.4	13.3	-0.1	15.1	97.4	19.6	-1.7	20.8	95.8
	(c)	17.7	-0.6	16.7	92.4	13.0	-1.4	15.1	97.6	21.6	-1.7	21.9	95.6
	(d)	16.8	-0.4	18.3	96.6	14.9	-0.3	18.6	98.4	20.7	-1.1	23.4	96.4
	(e)	23.0	-0.1	23.5	95.0	17.2	-0.8	18.1	96.0	28.3	-1.1	30.6	95.0
500	(a)	5.8	0.0	5.9	95.6	5.9	-0.1	5.9	95.4	5.9	-0.1	5.9	95.4
	(b)	10.9	-0.4	11.4	97.4	9.7	-0.4	9.5	93.8	12.5	-0.5	13.1	96.6
	(c)	11.0	-0.8	10.6	95.2	8.4	-0.8	9.0	96.4	12.9	-0.7	12.8	95.6
	(d)	10.3	0.5	11.3	96.8	9.6	-0.4	11.0	97.4	12.5	-0.6	13.6	97.4
	(e)	14.6	-1.0	14.9	93.8	10.4	0.4	11.2	97.2	15.9	-0.6	18.1	97.6
2000	(a)	2.8	0.3	2.9	96.2	2.6	0.3	2.9	96.6	2.6	0.3	2.9	96.6
	(b)	6.6	-0.2	6.6	94.0	5.2	0.0	5.3	95.4	7.5	-1.1	7.7	95.4
	(c)	6.2	-0.3	6.2	95.6	4.8	-0.1	4.7	95.2	7.6	-1.0	7.5	94.2
	(d)	5.9	0.0	6.2	96.2	5.1	0.1	5.5	96.2	6.8	-0.8	7.4	97.6
	(e)	7.2	-1.2	8.0	97.0	5.8	0.1	6.0	95.2	9.2	-1.1	9.7	96.8

Glossary of notation: p , the dimension of \mathbf{X} ; q , the sparsity level; n , the labeled data size; $m(\mathbf{X}) \equiv \mathbb{E}(Y | \mathbf{X})$; ESE, empirical standard error; ASE, average of estimated standard errors; CR, coverage rate of 95% confidence intervals.

Table 6: Covariates included for the data analysis in Section 5.

Variable name	Description
active	In your usual day, how active are you in 1971?
age	Age in 1971
alcoholfreq	How often do you drink in 1971?
allergies	Use allergies medication in 1971
asthma	DX asthma in 1971
cholesterol	Serum cholesterol (mg/100ml) in 1971
dbp	Diastolic blood pressure in 1982
education	Amount of education by 1971
exercise	In recreation, how much exercise in 1971?
ht	Height in centimeters in 1971
price71	Average tobacco price in state of residence 1971 (US\$2008)
price82	Average tobacco price in state of residence 1982 (US\$2008)
race	White, black or other in 1971
sbp	Systolic blood pressure in 1982
sex	Male or female
smokeintensity	Number of cigarettes smoked per day in 1971
smokeyrs	Years of smoking
tax71	Tobacco tax in state of residence 1971 (US\$2008)
tax82	Tobacco tax in state of residence 1971 (US\$2008)
wt71	Weight in kilograms in 1971

References

Andrews, D. W. (1995). Nonparametric kernel estimation for semiparametric models. *Econometric Theory*, pages 560–596.

Azriel, D., Brown, L. D., Sklar, M., Berk, R., Buja, A., and Zhao, L. (2022). Semi-supervised linear regression. *Journal of the American Statistical Association*, 117:2238–2251.

Belkin, M., Niyogi, P., and Sindhwani, V. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7(Nov):2399–2434.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

Bühlmann, P. and Van De Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Science & Business Media.

Chakrabortty, A. (2016). *Robust Semi-Parametric Inference in Semi-Supervised Settings*. PhD thesis, Harvard University, USA.

Chakrabortty, A. and Cai, T. (2018). Efficient and adaptive linear regression in semi-supervised settings. *Annals of Statistics*, 46(4):1541–1572.

Chakrabortty, A. and Dai, G. (2022). A general framework for treatment effect estimation in semi-supervised and high dimensional settings. *arXiv preprint arXiv:2201.00468*.

Chapelle, O., Schölkopf, B., and Zien, A. (2010). *Semi-Supervised Learning*. The MIT Press, 1st edition.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1):C1–C68.

Cozman, F. G. and Cohen, I. (2001). Unlabeled data can degrade classification performance of generative classifiers. Technical Report HPL-2001-234, HP Laboratories, Palo Alto, CA, USA.

Cozman, F. G., Cohen, I., and Cirelo, M. C. (2003). Semi-supervised learning of mixture models. In *Proceedings of the Twentieth ICML*, pages 99–106.

Ertefaie, A., Hejazi, N. S., and van der Laan, M. J. (2022). Nonparametric inverse probability weighted estimators based on the highly adaptive lasso. *Biometrics*.

Escanciano, J. C., Jacho-Chávez, D. T., and Lewbel, A. (2014). Uniform convergence of weighted sums of non and semiparametric residuals for estimation and testing. *Journal of Econometrics*, 178:426–443.

Flutre, T., Wen, X., Pritchard, J., and Stephens, M. (2013). A statistical framework for joint eQTL analysis in multiple tissues. *PLoS Genetics*, 9(5):e1003486.

Graham, B. S. (2011). Efficiency bounds for missing data models with semiparametric restrictions. *Econometrica*, 79(2):437–452.

Gronsbell, J. L. and Cai, T. (2018). Semi-supervised approaches to efficient evaluation of model prediction performance. *Journal of the Royal Statistical Society, Series B*, 80(3):579–594.

Hansen, B. E. (2008). Uniform convergence rates for kernel estimation with dependent data. *Econometric Theory*, 24:726–748.

Hernán, M. A. and Robins, J. M. (2020). *Causal Inference: What If*. Chapman & Hall/CRC.

Koenker, R. (2005). *Quantile Regression*. Cambridge University Press, Cambridge, UK.

Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327.

Lin, Q., Zhao, Z., and Liu, J. S. (2019). Sparse sliced inverse regression via lasso. *Journal of the American Statistical Association*, 114(528):1726–1739.

Little, R. J. and Rubin, D. B. (2019). *Statistical Analysis with Missing Data*, volume 793. John Wiley & Sons.

Liu, D., Lin, X., and Ghosh, D. (2007). Semiparametric regression of multidimensional genetic pathway data: least-squares kernel machines and linear mixed models. *Biometrics*, 63(4):1079–1088.

Mammen, E., Rothe, C., and Schienle, M. (2012). Nonparametric regression with nonparametrically generated covariates. *Annals of Statistics*, 40(2):1132–1170.

Mammen, E., Rothe, C., and Schienle, M. (2016). Semiparametric estimation with generated covariates. *Econometric Theory*, 32(5):1140–1177.

Masry, E. (1996). Multivariate local polynomial regression for time series: uniform strong consistency and rates. *Journal of Time Series Analysis*, 17(6):571–599.

McCaw, Z. R., Gaynor, S. M., Sun, R., and Lin, X. (2021). Cross-tissue eQTL mapping in the presence of missing data via surrogate outcome analysis. *bioRxiv*, pages 2020–11.

Michaelson, J. J., Loguercio, S., and Beyer, A. (2009). Detection and interpretation of expression quantitative trait loci (eQTL). *Methods*, 48(3):265–276.

Negahban, S. N., Ravikumar, P., Wainwright, M. J., and Yu, B. (2012). A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557.

Newey, W. K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. *Handbook of Econometrics*, 4:2111–2245.

Newey, W. K. and Robins, J. R. (2018). Cross-fitting and fast remainder rates for semiparametric estimation. *arXiv preprint arXiv:1801.09138*.

Nigam, K., McCallum, A. K., Thrun, S., and Mitchell, T. (2000). Text classification from labeled and unlabeled documents using EM. *Machine learning*, 39(2-3):103–134.

Seeger, M. (2002). Learning with labeled and unlabeled data. Technical report, Univ. Edinburgh, UK.

Tsiatis, A. (2007). *Semiparametric Theory and Missing Data*. Springer Science & Business Media.

Van der Vaart, A. W. (2000). *Asymptotic Statistics*, volume 3. Cambridge University Press.

Van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer.

Wainwright, M. J. (2019). *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press.

Zhang, A., Brown, L. D., and Cai, T. T. (2019). Semi-supervised inference: General theory and estimation of means. *Annals of Statistics*, 47(5):2538–2566.

Zhang, T. and Oles, F. (2000). The value of unlabeled data for classification problems. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 1191–1198.

Zhang, Y. and Bradic, J. (2022). High-dimensional semi-supervised learning: in search of optimal inference of the mean. *Biometrika*, 109:387–403.

Zhu, X. (2005). *Semi-Supervised Learning with Graphs*. PhD thesis, Carnegie Mellon University.

Zhu, X. and Goldberg, A. B. (2009). Introduction to semi-supervised learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 3(1):1–130.