

Semi-Supervised linear regression

David Azriel¹, Lawrence D. Brown², Michael Sklar³, Richard Berk², Andreas Buja² and Linda Zhao^{2*}

Technion – Israel institute of Technology¹, Wharton – University of Pennsylvania² and Stanford University³

Abstract

We study a regression problem where for some part of the data we observe both the *label* variable (Y) and the *predictors* (\mathbf{X}), while for other part of the data only the predictors are given. Such a problem arises, for example, when observations of the label variable are costly and may require a skilled human agent. If the conditional expectation $E[Y|\mathbf{X}]$ is exactly linear in \mathbf{X} then typically the additional observations of the \mathbf{X} 's do not contain useful information, but otherwise the unlabeled data can be informative. In this case, our aim is at constructing the best linear predictor. We suggest improved alternative estimates to the naive standard procedures that depend only on the labeled data. Our estimation method can be easily implemented and has simply described asymptotic properties. The new estimates asymptotically dominate the usual standard procedures under certain non-linearity condition of $E[Y|\mathbf{X}]$; otherwise, they are asymptotically equivalent. The performance of the new estimator for small sample size is investigated in an extensive simulation study. A real data example of inferring homeless population is used to illustrate the new methodology.

1 Introduction

1.1 Background and contribution

The term “semi-supervised learning” was coined in the machine learning literature to describe a situation in which some of the data is labeled while the rest of the data is unlabeled (Merz et al., 1992). Such situations occur when the label variable is difficult to observe and may require a complicated or expensive procedure. A typical example is web document classification, where the classification is done by a human agent while there are many more unlabeled on-line documents. Specifically, a sample of n observations from the joint distribution of (\mathbf{X}, Y) is given, where Y is a one-dimensional *label* variable and \mathbf{X} is a p -dimensional vector of *covariates* or *predictors*. Also, an additional sample of size m is observed where

*Corresponding author: David Azriel, davidazr@technion.ac.il. The work of Lawrence D. Brown and Linda Zhao was supported in part by NSF grant DMS-1512084. The work of Michael Sklar was supported in part by NSF grant NSF DMS RTG 1501767.

only \mathbf{X} is given. The purpose is to study procedures that make use of the additional unlabeled data to better capture the shape of the underlying joint distribution of the labeled data.

A large body of literature focuses on the case that Y takes a small number of values and the problem reduces to a classification task; see e.g., Zhu (2005) and references therein. When the predictors approximately lie in a low dimensional space, the unlabeled sample can be used to estimate the low dimensional structure (see e.g., Goldberg et al., 2009). Wang et al. (2009) divide the different methods into two approaches: distributional and margin-based. The distributional approach relies on an assumption relating the conditional expectation $E[Y|\mathbf{X}]$ to the marginal distribution of \mathbf{X} and the margin-based approach uses the extra information on \mathbf{X} for estimating the Bayes decision boundary; see Liang et al. (2008) for a Bayesian perspective within this basic approach. Other works consider continuous Y 's and use the unlabeled data to learn the structure of the \mathbf{X} 's in order to better estimate a non-parametric regression (Zhou, 2005; Lafferty and Wasserman, 2008; Johnson and Zhang, 2008). These works could be very helpful in situations where non-parametric regression is useful and unlabeled data are available.

Here we follow a different methodology. We aim to estimate the vector β composed of the parameters of the best linear predictor of Y , but we do not necessarily assume that $E[Y|\mathbf{X}]$ is exactly linear in \mathbf{X} . The methodology uses possible unmodeled non-linearities in $E(Y|\mathbf{X})$ and adapts information from the unlabeled data to provide an estimator for β . This estimator is asymptotically superior to the standard least squares estimate (LSE), both in terms of variance of β and mean squared error in predicting Y .

In the statistical literature, the typical approach to regression with unlabeled data may be best summarized by the following quote from Little (1992):

The related problem of missing values in the outcome Y was prominent in the early history of missing-data methods, but is less interesting in the following sense: If the X 's are complete and the missing values of Y are missing at random, then the incomplete cases contribute no information to the regression of Y on X_1, \dots, X_p .

But see Cochran (1977, Chapter 7) for a different view more closely in tune with our development.

Buja et al. (2017) show that β does not depend on the distribution of \mathbf{X} if and only if $E[Y|\mathbf{X}]$ is linear in \mathbf{X} . When the conditional expectation is linear, there is typically at most a limited amount of additional information in the unlabeled data. Brown (1990) shows that even in this case there may be some useful information, but it will not provide an asymptotic advantage in the manner we suggest in the current treatment when the conditional expectation is not linear. The assumption that $E[Y|\mathbf{X}]$ is linear is unrealistic in many situations, and we show that in the absence of such an assumption the unlabeled sample can be used to provide useful information for the estimation of β .

We consider two scenarios. In the first, the distribution of \mathbf{X} is known exactly. This is equivalent

to having infinitely many unlabeled observations ($m = \infty$). We call this the *full information* scenario. In the second scenario, which is more frequently encountered, $m < \infty$ but we assume that m is at least proportional to n . We call this *partial information*. In both situations we provide an asymptotically better estimate of β than the standard LSE, and hence also a better linear predictor. Our new procedure is closely related to the semi-supervised estimations of means, which is the topic of a recent manuscript (Zhang et al., 2017). This latter work constructs a semi-supervised regression estimator that is better than the sample mean. We build upon this work in two ways. First, we transform the regression problem into a mean estimation problem for each parameter via a “multiplication step”. Second, although it is proven by Zhang et al. that the semi-supervised estimate is superior to the sample mean, this does not directly imply superiority of our estimator to the LSE. For this purpose, a new argument is introduced based on an orthogonal decomposition of the error (see Section 1.2 below).

In a recent paper, Chakraborty and Cai (2018) also construct an estimator that is asymptotically superior to the LSE. Chakraborty and Cai estimate the non-linear part of $E(Y|\mathbf{X})$, non-parametrically, and then impute the missing Y ’s to aid estimation of β . This is challenging, and they suggest several methods for the nonparametric estimation.

Our approach, by contrast, squeezes extra efficiency from the LSE using minimal modeling. The two approaches are discussed further in Section 4, where we show that our estimator is the “simplest estimator” that dominates the LSE in some sense. We find that the suggested method is especially useful when the non-linear part is difficult to model.

1.2 The suggested methodology in a nutshell

To demonstrate our basic methodology, consider the simple linear regression model,

$$Y = \alpha + \beta X + \delta, \tag{1}$$

where α, β are least squares coefficients and δ is a remainder term; these will be defined in more detail shortly. When $E(X)$ and $E(X^2)$ are known, we construct an estimator of β that asymptotically dominates the least squares estimator. This is done by replacing model (1) with a different model where β is the intercept and also the expectation of the newly defined label variable. For this latter model, Zhang et al. (2017) show that the intercept estimator dominates the simple empirical estimate in a semi-supervised setting. Here we go one step further and show that the intercept estimator in this new model is better than the least squares estimator from the original model as explained below.

To present the intercept model and the methodology, begin by stating the general form of the best linear approximation. The approximation is used in two different models, and therefore we introduce now

the notation \mathbf{U} and W instead of \mathbf{X} and Y . Suppose that $W \in \mathbb{R}$ and $\mathbf{U} \in \mathbb{R}^d$ are random variables with joint distribution G and finite second moments. The best linear predictor is

$$\boldsymbol{\theta} = (\theta_1, \dots, \theta_d)^T = \arg \min_{\tilde{\boldsymbol{\theta}} \in \mathbb{R}^d} E(W - \sum_{j=1}^d U_j \tilde{\theta}_j)^2 = \{E(\mathbf{U}\mathbf{U}^T)\}^{-1} E(\mathbf{U}W). \quad (2)$$

Notice that this is a population version of the least squares where $\sum_{j=1}^d U_j \tilde{\theta}_j$ minimizes the L_2 distance in the population from W to any linear function of \mathbf{U} . It follows that

$$W = \sum_{j=1}^d \theta_j U_j + r, \quad (3)$$

where the remainder term $r = W - \sum_{j=1}^d \theta_j U_j$ is orthogonal to \mathbf{U} , i.e., $E(r\mathbf{U}) = 0$. Given a sample of n observations from G , the standard least squares estimate, $\hat{\boldsymbol{\theta}}_{LSE}$, satisfies asymptotically (White, 1980)

$$\sqrt{n}(\hat{\boldsymbol{\theta}}_{LSE} - \boldsymbol{\theta}) \xrightarrow{\mathcal{D}} N\left(0, \{E(\mathbf{U}\mathbf{U}^T)\}^{-1} E(r^2 \mathbf{U}\mathbf{U}^T) \{E(\mathbf{U}\mathbf{U}^T)\}^{-1}\right). \quad (4)$$

Unlike the standard fixed \mathbf{X} assumption, the asymptotic variance has a “sandwich” form, $\{E(\mathbf{U}\mathbf{U}^T)\}^{-1}$ forming the “bread” and $E(r^2 \mathbf{U}\mathbf{U}^T)$ the “meat”. See Buja et al. (2017) for further discussion of this form of the sandwich.

We now return to model (1) and assume that $E(X)$ and $E(X^2)$ are known. We can therefore assume w.l.o.g that X is standardized (i.e., $E(X) = 0$ and $E(X^2) = 1$). In this case we can write model (1) with $\alpha = E(Y)$, $\beta = E(XY)$ and $\delta = Y - (\alpha + \beta X)$; here we consider model (3) with $W = Y$, $U_1 = 1$, $U_2 = X$ and the remainder term is $r = \delta$. The standard LSE is a consistent estimate for β and satisfies, according to (4),

$$\sqrt{n}(\hat{\beta}_{LSE} - \beta) \xrightarrow{\mathcal{D}} N\{0, E(\delta X)^2\}.$$

Our aim is to better estimate $\beta = E(XY)$. To this end, we multiply (1) by X setting XY to be the labeled variable W . Furthermore, we also center the variable \mathbf{U} setting $E(XY) = \beta$ to be the intercept. Specifically,

$$XY = \alpha X + \beta X^2 + X\delta = \beta + aX + b(X^2 - 1) + \tilde{\delta}. \quad (5)$$

Now we consider model (3) with $W = XY$, $U_1 = 1$, $U_2 = X$, $U_3 = X^2 - 1$ and $r = \tilde{\delta}$; here a, b are θ_2, θ_3 defined by (2). In setting $\beta = E(XY)$ to be the intercept coefficient we used that $EX = E(X^2) - 1 = 0$, i.e., that we know the first and second moment of X . We define our new estimator, $\hat{\beta}_{TI}$ (TI for total information), to be the LSE intercept estimator of (5) based on n observations. The sandwich theorem (4) implies that

$$\sqrt{n}(\hat{\beta}_{TI} - \beta) \xrightarrow{\mathcal{D}} N(0, E\tilde{\delta}^2).$$

Since $\delta X - \tilde{\delta} = X(a - \alpha) + (X^2 - 1)(b - \beta)$ is a linear function of X, X^2 , then it is orthogonal to the remainder term of (5), which is $\tilde{\delta}$, (i.e., $E(\delta X - \tilde{\delta})\tilde{\delta} = 0$) and therefore,

$$E(\delta X)^2 = E\{\tilde{\delta} + (\delta X - \tilde{\delta})\}^2 = E\tilde{\delta}^2 + E(\delta X - \tilde{\delta})^2. \quad (6)$$

This implies that the asymptotic variance of $\hat{\beta}_{TI}$ is smaller than that of $\hat{\beta}_{LSE}$ with equality iff $\delta X \equiv \tilde{\delta}$. The latter occurs iff $\alpha = a$ and $b = \beta$ or equivalently that δ is uncorrelated with both X^2 and X^3 . This occurs when the non-linear part of $E[Y|X]$ (if exists) is uncorrelated with X^2 and X^3 . In this case, Models (1) and (5) are essentially the same model and nothing is gained in the new methodology. On the other hand, when δ is correlated with either X^2 or X^3 , then Models (1) and (5) are different and $\hat{\beta}_{TI}$ has smaller asymptotic variance than $\hat{\beta}_{LSE}$. We further show that a similar decomposition to (6) holds for the partial information case and we generalize the results to the p -dimensional case. We also show that smaller variance implies smaller prediction error and therefore $\hat{\beta}_{TI}$ provides better predictions than $\hat{\beta}_{LSE}$.

The rest of the paper is organized as follows. Section 2 provides the basic setting and the loss functions that we use. The main results are given in Section 3 and in Section 4 we discuss the relation between our method and semi-parametric efficiency. Section 5 describes an extensive simulation study and Section 6 discusses estimation of the asymptotic covariance matrix of the estimates. An implementation of the new methodology to infer homeless population in Los Angeles is discussed in Section 7. Section 8 concludes with final remarks. The proofs are given in Section 9.

2 Preliminaries

Consider a sample of n i.i.d observations $(\mathbf{X}^{(1)}, Y^{(1)}), \dots, (\mathbf{X}^{(n)}, Y^{(n)})$ from a joint distribution G , where $\mathbf{X} \in \mathbb{R}^p$, $Y \in \mathbb{R}$, and an additional set of m independent observations $(\mathbf{X}^{(n+1)}, \dots, \mathbf{X}^{(n+m)})$ from the marginal distribution of \mathbf{X} . We use super-index to denote the number of the observation, and sub-index to denote coordinates of \mathbf{X} . The notation \mathbf{X}, Y without super-index denotes a random vector whose distribution is G .

We write $\vec{\mathbf{X}} = (1, X_1, \dots, X_p)^T$ to be a vector \mathbf{X} with an additional constant 1 to accommodate an intercept term. Assume that the second moments of G exists and that the matrix $E(\vec{\mathbf{X}}\vec{\mathbf{X}}^T)$ is invertible. Then, we can define

$$(\alpha, \beta) = \arg \min_{\tilde{\alpha} \in \mathbb{R}, \tilde{\beta} \in \mathbb{R}^p} E(Y - \tilde{\alpha} - \tilde{\beta}^T X)^2 = \left\{ E(\vec{\mathbf{X}}\vec{\mathbf{X}}^T) \right\}^{-1} E(\vec{\mathbf{X}}Y). \quad (7)$$

In the presence of non-linearity, β is still a meaningful parameter that describes the overall association between Y and \mathbf{X} (Buja et al., 2017). We have in mind two related purposes. The first purpose is just

to better estimate the parameters of interest, while the second purpose pertains to prediction. The latter is formalized now. Suppose that an independent observation $(\mathbf{X}^*, Y^*) \sim G$ is given. The optimal linear predictor is $\alpha + \boldsymbol{\beta}^T X^*$. We consider the excess loss of an estimator $\tilde{\alpha}, \tilde{\boldsymbol{\beta}}$

$$L(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) = \left(Y^* - \tilde{\alpha} - \tilde{\boldsymbol{\beta}}^T \mathbf{X}^* \right)^2 - \left(Y^* - \alpha - \boldsymbol{\beta}^T \mathbf{X}^* \right)^2.$$

We have that (see Lemma 1)

$$EL(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) = E \left[\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{ \bar{X}_j - E(X_j) \} - E(Y) \right]^2 + E \left\{ (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{M} (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}}) \right\} = E\tilde{L}(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}),$$

where \tilde{L} is the expression inside the expectation and \mathbf{M} is the covariance matrix of \mathbf{X} . Notice that L and \tilde{L} have the same expectation but they are different random variables. We further define the asymptotic risk as

$$R(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) = \lim_{B \rightarrow \infty} \lim_{n \rightarrow \infty} E \min \{ n\tilde{L}(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}), B \}.$$

The loss is of order $1/n$ and therefore we consider expectation of $n\tilde{L}$. This is truncated by an arbitrarily large number B , since when the loss is large, it makes sense not to penalize any further. Also, the truncation helps to avoid issues of uniform integrability. This is done for example in Le Cam and Yang (1990), Chapter 5. The following proposition provides a simple expression for the asymptotic risk.

Proposition 1. *Let $\tilde{\alpha}, \tilde{\boldsymbol{\beta}}$ satisfies*

$$\tilde{\alpha} = \bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \bar{X}_j, \quad (8)$$

and assume that $\tilde{\boldsymbol{\beta}}$ satisfies $\sqrt{n}(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{\mathcal{D}} N(0, \boldsymbol{\Sigma})$ then,

$$R(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) = E \left[Y - \sum_{j=1}^p \beta_j \{ X_j - E(X_j) \} - E(Y) \right]^2 + \text{Trace}(\mathbf{M}\boldsymbol{\Sigma}), \quad (9)$$

where \mathbf{M} is the covariance matrix of \mathbf{X} .

The first term in (9) does not depend on the distribution of $\tilde{\boldsymbol{\beta}}$. Hence, Proposition 1 shows that the excess risk is minimized when $\text{Trace}(\mathbf{M}\boldsymbol{\Sigma})$ is small. Thus, we aim at estimators $\tilde{\boldsymbol{\beta}}$ with asymptotic distribution $N(0, \boldsymbol{\Sigma})$ such that $\boldsymbol{\Sigma}$ is “smaller” than the covariance matrix of LSE, in the sense that the difference is positive semidefinite. Such an estimator asymptotically better estimates $\boldsymbol{\beta}$ and also has smaller asymptotic excess risk.

For two estimates $(\tilde{\alpha}^{(1)}, \tilde{\boldsymbol{\beta}}^{(1)})$ and $(\tilde{\alpha}^{(2)}, \tilde{\boldsymbol{\beta}}^{(2)})$, Proposition 1 implies that

$$EL(\tilde{\alpha}^{(1)}, \tilde{\boldsymbol{\beta}}^{(1)}) - EL(\tilde{\alpha}^{(2)}, \tilde{\boldsymbol{\beta}}^{(2)}) \approx \text{Trace} \left\{ \mathbf{M}(\boldsymbol{\Sigma}^{(1)} - \boldsymbol{\Sigma}^{(2)}) \right\} / n,$$

and therefore the difference of the prediction errors is

$$E \left\{ Y^* - \tilde{\alpha}^{(1)} - \left(\tilde{\boldsymbol{\beta}}^{(1)} \right)^T \mathbf{X}^* \right\}^2 - E \left\{ Y^* - \tilde{\alpha}^{(2)} - \left(\tilde{\boldsymbol{\beta}}^{(2)} \right)^T \mathbf{X}^* \right\}^2 \approx \text{Trace} \left\{ \mathbf{M}(\boldsymbol{\Sigma}^{(1)} - \boldsymbol{\Sigma}^{(2)}) \right\} / n. \quad (10)$$

It follows that the difference of the errors of the prediction of the mean is also

$$E \left\{ E(Y^*|\mathbf{X}^*) - \tilde{\alpha}^{(1)} - \left(\tilde{\boldsymbol{\beta}}^{(1)} \right)^T \mathbf{X}^* \right\}^2 - E \left\{ E(Y^*|\mathbf{X}^*) - \tilde{\alpha}^{(2)} - \left(\tilde{\boldsymbol{\beta}}^{(2)} \right)^T \mathbf{X}^* \right\}^2 \approx \text{Trace} \left\{ \mathbf{M}(\boldsymbol{\Sigma}^{(1)} - \boldsymbol{\Sigma}^{(2)}) \right\} / n. \quad (11)$$

Notice that we estimate the intercept, as in (8), using \bar{X}_j rather than the possibly known expectations $E(X_j)$. This is in accordance with the findings of Zhang et al. (2017).

3 Main results

In this section we provide the main theoretical results of the paper. Before we deal with the p-dimensional case, we first introduce the results for one-dimensional \mathbf{X} . The reason is twofold: first, the presentation in the one-dimensional case is simpler and captures the main ideas, and second, our results for the p-dimensional \mathbf{X} are obtained by reducing the regression problem to p problems, each of which is closely related to the one dimensional \mathbf{X} regression.

3.1 One dimensional X

In this section we study the one dimensional case. Summary of the notation used here is presented in Table 1. When G has finite second moments we can write

$$Y = \alpha + \beta X + \delta, \quad (12)$$

where $\alpha = E(Y) - \beta E(X)$, $\beta = \frac{\text{Cov}(X,Y)}{\text{Var}(X)}$ and $\delta = Y - \alpha - \beta X$. Equation (12) is the best linear approximation in the population in the sense that $\alpha + \beta X$ minimizes $E(\tilde{\alpha} + \tilde{\beta}X - Y)^2$ over all $\tilde{\alpha}, \tilde{\beta} \in \mathbb{R}$ and δ is orthogonal to $1, X$, i.e., $E(\delta) = E(\delta X) = 0$. The regular LSE is

$$\hat{\beta}_{LSE} = \frac{\sum_{i=1}^n (X^{(i)} - \bar{X})Y^{(i)}}{\sum_{i=1}^n (X^{(i)} - \bar{X})^2}.$$

This is a consistent and asymptotically unbiased and asymptotically normal estimator for β . See Buja et al. (2017) for contemporary discussion.

For the total information estimator (TI) we consider the following regression model, which is obtained by multiplying (12) by $\frac{X-E(X)}{\text{Var}(X)}$,

$$W = \beta + aU_1 + bU_2 + \tilde{\delta}, \quad (13)$$

where $W = \frac{Y\{X-E(X)\}}{\text{Var}(X)}$, $U_1 = \frac{X-E(X)}{\text{Var}(X)}$, $U_2 = \frac{\{X-E(X)\}X}{\text{Var}(X)} - 1$, a, b are the coefficients θ_2, θ_3 of the best linear approximation defined by (2) and $\tilde{\delta}$ is the remainder term. Here again $\beta + aU_1 + bU_2$ is the best linear approximation of W in the population. The multiplication term makes the expectation of W to be $E \left[\frac{Y\{X-E(X)\}}{\text{Var}(X)} \right] = \frac{\text{Cov}(X,Y)}{\text{Var}(X)} = \beta$. The covariates U_1, U_2 are the predictors $1, X$ in (12) multiplied by

Table 1: Summary of the notation used in Section 3.1.

Basic model	$Y = \alpha + \beta X + \delta$	$\alpha = E(Y) - \beta E(X)$, $\beta = \frac{\text{Cov}(X,Y)}{\text{Var}(X)}$ and $\delta = Y - \alpha - \beta X$
Intercept model	$W = \beta + aU_1 + bU_2 + \tilde{\delta}$	$W = \frac{Y\{X-E(X)\}}{\text{Var}(X)}$, $U_1 = \frac{X-E(X)}{\text{Var}(X)}$, $U_2 = \frac{\{X-E(X)\}X}{\text{Var}(X)} - 1$ a, b are θ_2, θ_3 defined by (2) for model (13) and $\tilde{\delta}$ is the remainder term
TI estimator	$\hat{\beta}_{TI} = \bar{W} - \hat{a}\bar{U}_1 - \hat{b}\bar{U}_2$	$\bar{W} = \frac{1}{n} \sum_{i=1}^n W^{(i)}$, $\bar{U}_1 = \frac{1}{n} \sum_{i=1}^n U_1^{(i)}$, $\bar{U}_2 = \frac{1}{n} \sum_{i=1}^n U_2^{(i)}$ \hat{a}, \hat{b} are the LSE of model (13)
PI estimator	$\hat{\beta}_{PI} = \bar{\check{W}} - \hat{a}\bar{\check{U}}_1 - \hat{b}\bar{\check{U}}_2$	$\bar{\check{W}}, \bar{\check{U}}_1, \bar{\check{U}}_2$ are the means over the labeled sample, $\check{W}^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}Y^{(i)}}{\text{Var}(X)}$, $\check{U}_1^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}}{\text{Var}(X)}$ $\check{U}_2^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}X^{(i)}}{\text{Var}(X)} - 1$, \hat{a}, \hat{b} are the LSE of model (13) with $\check{W}, \check{U}_1, \check{U}_2$ instead of W, U_1, U_2
Asymptotics	$\nu = \lim \frac{n}{n+m}$ $\sigma_{LSE}^2, \sigma_{PI}^2, \sigma_{TI}^2$	n (m) the labeled (unlabeled) sample size asymptotic variance of the LSE, PI, TI estimators as in Theorem 1
Main results	$\sigma_{LSE}^2 = \sigma_{TI}^2 + \sigma_{\text{diff}}^2$ $\sigma_{PI}^2 = \sigma_{TI}^2 + \nu \sigma_{\text{diff}}^2$	$\sigma_{\text{diff}}^2 = E \left[\frac{\delta\{X-E(X)\}}{\text{Var}(X)} - \tilde{\delta} \right]^2$

$\frac{X-E(X)}{\text{Var}(X)}$. We subtract 1 in U_2 in order to set $E(U_2) = 0$. Since $E(U_1) = E(U_2) = 0$, then the intercept is $E(W) = \beta$. Thus, we define the total information estimator to be the intercept estimator of (13), i.e.,

$$\hat{\beta}_{TI} = \bar{W} - \hat{a}\bar{U}_1 - \hat{b}\bar{U}_2, \quad (14)$$

where \hat{a}, \hat{b} are the regular LSE of model (13), and $\bar{\cdot}$ denotes the mean over the supervised sample with n observations.

For the partial information (PI) estimator we estimate W, U_1, U_2 as follows:

$$\check{W}^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}Y^{(i)}}{\check{\text{Var}}(X)}, \check{U}_1^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}}{\check{\text{Var}}(X)} \text{ and } \check{U}_2^{(i)} = \frac{\{X^{(i)} - \check{E}(X)\}X^{(i)}}{\check{\text{Var}}(X)} - 1, i = 1, \dots, n,$$

where $\check{E}(X) = \frac{1}{n+m} \sum_{i=1}^{n+m} X^{(i)}$, $\check{\text{Var}}(X) = \check{E}(X^2) - \{\check{E}(X)\}^2$, and $\check{E}(X^2) = \frac{1}{n+m} \sum_{i=1}^{n+m} \{X^{(i)}\}^2$. The partial information estimator is

$$\hat{\beta}_{PI} = \bar{\check{W}} - \hat{a}\bar{\check{U}}_1 - \hat{b}\bar{\check{U}}_2,$$

where \hat{a}, \hat{b} are the regular LSE of the regression model (13) with $\check{W}, \check{U}_1, \check{U}_2$ instead of W, U_1, U_2 . (The estimates \hat{a}, \hat{b} are different in the total and partial information cases but the same notation is presented for simplicity.) We use $\bar{\cdot}$ (respectively, $\check{\cdot}$) to denote empirical mean with respect to the labeled n (respectively,

full $n + m$) sample. The following theorem states the asymptotic distribution of $\hat{\beta}_{LSE}$, $\hat{\beta}_{TI}$ and $\hat{\beta}_{PI}$. The first part of the theorem is known and is stated here for comparison purposes.

Theorem 1. (i) Suppose that $\text{Var}(X) \in (0, \infty)$ and that $\delta\{X - E(X)\}$ has finite second moment, then,

$$\sqrt{n}(\hat{\beta}_{LSE} - \beta) \xrightarrow{\mathcal{D}} N(0, \sigma_{LSE}^2),$$

$$\text{where } \sigma_{LSE}^2 = E \left[\frac{\delta^2 \{X - E(X)\}^2}{\{\text{Var}(X)\}^2} \right].$$

(ii) Suppose further that the vector (W, U_1, U_2) has finite second moments and that the matrix $E(\mathbf{U}\mathbf{U}^T)$ is invertible for $\mathbf{U} = (1, U_1, U_2)^T$. Then,

$$\sqrt{n}(\hat{\beta}_{TI} - \beta) \xrightarrow{\mathcal{D}} N(0, \sigma_{TI}^2),$$

$$\text{where } \sigma_{TI}^2 = E(\tilde{\delta}^2) \text{ and } \sigma_{LSE}^2 = \sigma_{TI}^2 + \sigma_{\text{diff}}^2 \text{ where } \sigma_{\text{diff}}^2 = E \left[\frac{\delta\{X - E(X)\}}{\text{Var}(X)} - \tilde{\delta} \right]^2.$$

(iii) Suppose further that $\lim \frac{n}{n+m} = \nu$, then,

$$\sqrt{n}(\hat{\beta}_{PI} - \beta) \xrightarrow{\mathcal{D}} N(0, \sigma_{PI}^2),$$

$$\text{where } \sigma_{PI}^2 = \sigma_{TI}^2 + \nu \sigma_{\text{diff}}^2.$$

Therefore, if $\sigma_{\text{diff}}^2 > 0$ then $\sigma_{TI}^2 < \sigma_{LSE}^2$ and if further $\nu < 1$ then $\sigma_{PI}^2 < \sigma_{LSE}^2$.

Corollary 1. Theorem 1 and Proposition 1 imply that $R(\hat{\alpha}_{LSE}, \hat{\beta}_{LSE}) - R(\hat{\alpha}_{PI}, \hat{\beta}_{PI}) = (1 - \nu) \sigma_{\text{diff}}^2 \text{Var}(X)$.

The improvement of $\hat{\beta}_{TI}$ and $\hat{\beta}_{PI}$ over $\hat{\beta}_{LSE}$ depends on the assumption that $\sigma_{\text{diff}}^2 > 0$. The quantity σ_{diff}^2 measures the difference between the original regression model (12) and the intercept model (13). If $E(\delta X^2) = E(\delta X^3) = 0$ then $a = \alpha$ and $b = \beta$ and the two models are essentially the same, in which case $\sigma_{\text{diff}}^2 = 0$, otherwise $\sigma_{\text{diff}}^2 > 0$. In other words, when $E[Y|X]$ is non-linear and the non-linear part is correlated with either X^2 or X^3 then $\sigma_{\text{diff}}^2 > 0$.

3.2 Multidimensional X

We now consider the general p -dimensional case as described in the beginning of Section 2. The notation of this section is summarized in Table 2. The model can be written as

$$Y = \alpha + \beta_1 X_1 + \cdots + \beta_p X_p + \delta, \tag{15}$$

where α, β are the coefficients of the best linear predictor in the population defined by (7). The remainder term $\delta = Y - \sum_{j=1}^p \beta_j X_j$ satisfies $E(\delta) = E(\delta X_1) = \cdots = E(\delta X_p) = 0$. Our aim here is at estimating β . We use the adjustment representation of Buja et al. (2017) that reduces the p -dimensional estimation

Table 2: Summary of the notation used in Section 3.2 when j is the chosen coordinate.

Basic model	$Y = \alpha + \sum_{j=1}^p \beta_j X_j + \delta$	$\alpha, \boldsymbol{\beta}$ are defined by (7)
Adjusted regressor	$X_{j\bullet} = X_j - \vec{\mathbf{X}}_{-j}^T \boldsymbol{\beta}_{-j\bullet}$	$\boldsymbol{\beta}_{-j\bullet} = \left\{ E \left(\vec{\mathbf{X}}_{-j} \vec{\mathbf{X}}_{-j}^T \right) \right\}^{-1} E \left(\vec{\mathbf{X}}_{-j} X_j \right)$ $\vec{\mathbf{X}}_{-j} = (1, X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_p)^T$
Intercept model	$W_j = \beta_j + aU_1 + \sum_{j'=1}^p b_{j'} U_{j'+1} + \tilde{\delta}_j$	$W_j = \frac{Y X_{j\bullet}}{E(X_{j\bullet}^2)}, U_1 = \frac{X_{j\bullet}}{E(X_{j\bullet}^2)}$ $U_{j'+1} = \frac{X_{j'} X_{j\bullet}}{E(X_{j\bullet}^2)}$ for $j' \neq j$, $U_{j+1} = \frac{X_j X_{j\bullet}}{E(X_{j\bullet}^2)} - 1$, a, \mathbf{b} are $\theta_2, \dots, \theta_{p+1}$ defined by (2) for model (17), $\tilde{\delta}_j$ is the remainder term
TI estimator	$\{\hat{\boldsymbol{\beta}}_{TI}\}_j = \bar{W}_j - \hat{a}\bar{U}_1 - \sum_{j'=1}^p \hat{b}_{j'} \bar{U}_{j'+1}$	$\bar{\cdot}$ is the mean over the labeled sample $\hat{a}, \hat{\mathbf{b}}$ are the LSE of model (17)
PI estimator	$\{\hat{\boldsymbol{\beta}}_{PI}\}_j = \check{W}_j - \hat{a}\check{U}_1 - \sum_{j'=1}^p \hat{b}_{j'} \check{U}_{j'+1}$	$\bar{\cdot}$ is the mean over the labeled sample $\check{W}_j = \frac{Y X_{j\bullet}}{E(X_{j\bullet}^2)}, \check{U}_1 = \frac{X_{j\bullet}}{E(X_{j\bullet}^2)}$ $\check{U}_{j'+1} = \frac{X_{j'} X_{j\bullet}}{E(X_{j\bullet}^2)}$ for $j' \neq j$, $\check{U}_{j+1} = \frac{X_j X_{j\bullet}}{E(X_{j\bullet}^2)} - 1$ $\hat{a}, \hat{\mathbf{b}}$ are the LSE of model (17) with $(\check{W}_j, \check{U}_1, \dots, \check{U}_{p+1})$ replacing $(W_j, U_1, \dots, U_{p+1})$
Asymptotics	$\nu = \lim \frac{n}{n+m}$ $\boldsymbol{\Sigma}_{LSE}, \boldsymbol{\Sigma}_{PI}, \boldsymbol{\Sigma}_{TI}$	n (m) the labeled (unlabeled) sample size asymptotic covariance matrix of LSE, PI, TI as in Theorem 2
Main results	$\boldsymbol{\Sigma}_{LSE} = \boldsymbol{\Sigma}_{TI} + \boldsymbol{\Sigma}_{\text{diff}}$ $\boldsymbol{\Sigma}_{PI} = \boldsymbol{\Sigma}_{TI} + \nu \boldsymbol{\Sigma}_{\text{diff}}$	$\boldsymbol{\Sigma}_{\text{diff}} = \text{Cov}(\delta \mathbf{X}_{\bullet} - \tilde{\delta})$

procedure to p separate simple regression problems. Correspondingly, we will define our new estimates by solving p mean-estimation-problems separately, one for each coordinate j .

Let $\vec{\mathbf{X}}_{-j} = (1, X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_p)^T$ and let

$$\boldsymbol{\beta}_{-j\bullet} = \arg \min_{\boldsymbol{\beta}} E \left(X_j - \boldsymbol{\beta}^T \vec{\mathbf{X}}_{-j} \right)^2 = \left\{ E \left(\vec{\mathbf{X}}_{-j} \vec{\mathbf{X}}_{-j}^T \right) \right\}^{-1} E \left(\vec{\mathbf{X}}_{-j} X_j \right). \quad (16)$$

Now, define $X_{j\bullet} = X_j - \vec{\mathbf{X}}_{-j}^T \boldsymbol{\beta}_{-j\bullet}$. Each β_j can be written in the one dimensional form $\beta_j = \frac{E(Y X_{j\bullet})}{E(X_{j\bullet}^2)}$.

The standard LSE can be viewed in a similar manner. Let $\mathbf{Y} = (Y^{(1)}, \dots, Y^{(n)})^T$, $\mathbf{X}_j = (X_j^{(1)}, \dots, X_j^{(n)})^T$ and let $\mathbf{X}_{-j} = (\mathbf{1}, \mathbf{X}_1, \dots, \mathbf{X}_{j-1}, \mathbf{X}_{j+1}, \dots, \mathbf{X}_p)$. Define also $\hat{\boldsymbol{\beta}}_{-j\bullet} = \left\{ \mathbf{X}_{-j}^T \mathbf{X}_{-j} \right\}^{-1} \mathbf{X}_{-j}^T \mathbf{X}_j$ and $\mathbf{X}_{j\bullet} = \mathbf{X}_j - \mathbf{X}_{-j} \hat{\boldsymbol{\beta}}_{-j\bullet}$; then, $\{\hat{\boldsymbol{\beta}}_{LSE}\}_j = \frac{\langle \mathbf{Y}, \mathbf{X}_{j\bullet} \rangle}{\|\mathbf{X}_{j\bullet}\|^2}$. Recall that $\bar{\cdot}$ denotes mean over the labeled sample; thus, $\bar{\cdot}$

denotes adjustments over the labeled sample, whereas \bullet denotes adjustments over the population.

The total information estimator is the intercept estimator of the regression model in the population obtained by multiplying (15) by $\frac{X_{j\bullet}}{E(X_{j\bullet}^2)}$, that is,

$$W_j = \beta_j + aU_1 + \sum_{j'=1}^p b_{j'}U_{j'+1} + \tilde{\delta}_j, \quad (17)$$

where

$$W_j = \frac{YX_{j\bullet}}{E(X_{j\bullet}^2)}, \quad U_1 = \frac{X_{j\bullet}}{E(X_{j\bullet}^2)}, \quad U_{j'+1} = \frac{X_{j'}X_{j\bullet}}{E(X_{j\bullet}^2)} \text{ for } j' \neq j, \text{ and } U_{j+1} = \frac{X_jX_{j\bullet}}{E(X_{j\bullet}^2)} - 1 \quad (18)$$

are $Y, 1, X_1, \dots, X_p$ multiplied by $\frac{X_{j\bullet}}{E(X_{j\bullet}^2)}$ and a, \mathbf{b} are $\theta_2, \dots, \theta_{p+1}$ defined by (2) for $W = W_j$ and $\mathbf{U} = (U_1, \dots, U_{p+1})$. Since $E(U_1) = \dots = E(U_{p+1}) = 0$, then the intercept is $E(W_j) = \beta_j$. In setting $E(W_j) = \beta_j$ and also the expectations of the U 's to be zero, we exploited the knowledge of the moments of \mathbf{X} , yielding higher efficiency as shown below. The remainder term in (17), $\tilde{\delta}_j = W_j - \beta_j - aU_1 - \sum_{j'=1}^p b_{j'}U_{j'+1}$ is orthogonal to (U_1, \dots, U_{p+1}) . Notice that the vector (U_1, \dots, U_{p+1}) depends on j but this is suppressed in the notation.

Specifically, we define the total information (TI) estimator to be

$$\{\hat{\beta}_{TI}\}_j = \bar{W}_j - \hat{a}\bar{U}_1 - \sum_{j'=1}^p \hat{b}_{j'}\bar{U}_{j'+1},$$

where $\hat{a}, \hat{\mathbf{b}}$ are the regular LSE of (17), and $\bar{\cdot}$ denotes the mean over the labeled sample with n observations.

To define the partial information estimator let

$$\beta_{-j\check{\bullet}} = \left\{ \check{E} \left(\vec{\mathbf{X}}_{-j} \vec{\mathbf{X}}_{-j}^T \right) \right\}^{-1} \check{E} \left(\vec{\mathbf{X}}_{-j} X_j \right), \quad (19)$$

where \check{E} is the empirical mean based on the full \mathbf{X} sample of size $n + m$. Now define $X_{j\check{\bullet}} = X_j - \left\{ \vec{\mathbf{X}}_{-j} \right\}^T \beta_{-j\check{\bullet}}$ and

$$\check{W}_j = \frac{YX_{j\check{\bullet}}}{\check{E}(X_{j\check{\bullet}}^2)}, \quad \check{U}_1 = \frac{X_{j\check{\bullet}}}{\check{E}(X_{j\check{\bullet}}^2)}, \quad \check{U}_{j'+1} = \frac{X_{j'}X_{j\check{\bullet}}}{\check{E}(X_{j\check{\bullet}}^2)} \text{ for } j' \neq j, \text{ and } \check{U}_{j+1} = \frac{X_jX_{j\check{\bullet}}}{\check{E}(X_{j\check{\bullet}}^2)}.$$

Here $\check{\bullet}$ denotes adjustments over the full $n + m$ sample. The partial information (PI) estimator is

$$\{\hat{\beta}_{PI}\}_j = \check{W}_j - \hat{a}\check{U}_1 - \sum_{j'=1}^p \hat{b}_{j'}\check{U}_{j'+1},$$

where $\hat{a}, \hat{b}_{j'}$ are the regular LSE of model (17) with $(\check{W}_j, \check{U}_1, \dots, \check{U}_{p+1})$ replacing $(W_j, U_1, \dots, U_{p+1})$.

The following theorem establishes the asymptotic distribution of the estimates and states conditions under which $\hat{\beta}_{TI}$ and $\hat{\beta}_{PI}$, asymptotically dominate $\hat{\beta}_{LSE}$. The asymptotic distribution of $\hat{\beta}_{LSE}$ is already known (White, 1980) and it is presented here so that the comparison to $\hat{\beta}_{TI}$ and $\hat{\beta}_{PI}$ can be made.

Theorem 2. (i) Suppose that the vector $\mathbf{X}_\bullet = \left\{ \frac{X_{1\bullet}}{E(X_{1\bullet}^2)}, \dots, \frac{X_{p\bullet}}{E(X_{p\bullet}^2)} \right\}$ is well defined (i.e., the projections in (16) exist and $E(X_{j\bullet})^2$ is positive and finite for $j = 1, \dots, p$) and that $\delta\mathbf{X}_\bullet$ has finite second moments, then,

$$\sqrt{n}(\hat{\beta}_{LSE} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{LSE}),$$

where $\Sigma_{LSE} = Cov(\delta\mathbf{X}_\bullet)$.

(ii) Suppose further that for each $j = 1, \dots, p$, the vector $(W_j, U_1, U_2, \dots, U_{p+1})$ has finite second moments and that the matrix $E(\mathbf{U}_j \mathbf{U}_j^T)$ is invertible for $\mathbf{U}_j = (1, U_1, U_2, \dots, U_{p+1})^T$; then,

$$\sqrt{n}(\hat{\beta}_{TI} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{TI}),$$

where $\Sigma_{TI} = Cov(\tilde{\delta})$ for $\tilde{\delta} = (\tilde{\delta}_1, \dots, \tilde{\delta}_p)$. Furthermore, we have that $\Sigma_{LSE} = \Sigma_{TI} + \Sigma_{\text{diff}}$ where $\Sigma_{\text{diff}} = Cov(\mathbf{V})$ for $\mathbf{V} = \delta\mathbf{X}_\bullet - \tilde{\delta}$.

(iii) Suppose further that $\lim_{n+m} \frac{n}{n+m} = \nu$, then,

$$\sqrt{n}(\hat{\beta}_{PI} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{PI}),$$

where $\Sigma_{PI} = \Sigma_{TI} + \nu\Sigma_{\text{diff}}$.

Therefore, if Σ_{diff} is not the zero matrix, then $\Sigma_{LSE} - \Sigma_{TI}$ is positive definite, and if further $\nu < 1$, then $\Sigma_{LSE} - \Sigma_{PI}$ is also positive definite.

Corollary 2. Theorem 2 and Proposition 1 imply that $R(\hat{\alpha}_{LSE}, \hat{\beta}_{LSE}) - R(\hat{\alpha}_{PI}, \hat{\beta}_{PI}) = \text{Trace}\{(1 - \nu)\Sigma_{\text{diff}}\mathbf{M}\}$, where $\mathbf{M} = Cov(\mathbf{X})$.

In short, Theorem 2 states that if the regression models (15) and (17) are well defined and the residuals have finite second moments, then $\Sigma_{LSE} - \Sigma_{TI}$ is positive semi-definite and it is strictly positive unless $\mathbf{V} \equiv 0$. If further, the unlabeled sample size is not negligible, i.e. $\lim m/n > 0$, then $\Sigma_{LSE} - \Sigma_{PI}$ is also positive definite.

As in the one-dimensional case, $\hat{\beta}_{TI}$ and $\hat{\beta}_{PI}$ improve over $\hat{\beta}_{LSE}$ only when Σ_{diff} is not the zero matrix or equivalently that \mathbf{V} is not zero. For each j , $Var(V_j)$ measures the difference between the original model (15) and the intercept model (17). When

$$E(X_{j\bullet}^2 \delta) = E(X_{j\bullet}^2 X_{1\bullet} \delta) = \dots = E(X_{j\bullet}^2 X_{p\bullet} \delta) = 0 \quad (20)$$

then $\delta X_{j\bullet} \equiv \delta_j$ and the two models are essentially the same, in which case $V_j \equiv 0$. Otherwise, if for some j (20) does not hold, then $\hat{\beta}_{PI}$ improves over $\hat{\beta}_{LSE}$ when $\nu < 1$. In the one-dimensional case, (20) is equivalent to $E(X^2 \delta) = E(X^3 \delta) = 0$, i.e., that X^2 and X^3 are uncorrelated with δ . Generally, (20) implies that certain linear combinations of $E(X_j X_{j'} \delta)$ and $E(X_j X_{j'} X_{j''} \delta)$ are 0. Thus, roughly speaking, when $E[Y|\mathbf{X}]$ is non-linear and the non-linear part is correlated with several second or third moments of \mathbf{X} , then we expect improvement of $\hat{\beta}_{PI}$.

3.3 New methodology (summary)

We now provide a step-by-step description of the new methodology for estimation of β .

For each $j = 1, \dots, p$:

1. (Adjust the regressors) Let $\beta_{-j\check{\bullet}}$ as defined in (19) and define $X_{j\check{\bullet}}^{(i)} = X_j^{(i)} - \{\bar{\mathbf{X}}_{-j}^{(i)}\}^T \beta_{-j\check{\bullet}}$ for $i = 1, \dots, n$.
2. (Define the intercept model) Define for $i = 1, \dots, n$

$$\check{W}_j^{(i)} = \frac{Y^{(i)} X_{j\check{\bullet}}^{(i)}}{\check{E}(X_{j\check{\bullet}}^2)}, \check{U}_1^{(i)} = \frac{X_{j\check{\bullet}}^{(i)}}{\check{E}(X_{j\check{\bullet}}^2)}, \check{U}_{j'+1}^{(i)} = \frac{X_{j'\check{\bullet}}^{(i)} X_{j\check{\bullet}}^{(i)}}{\check{E}(X_{j\check{\bullet}}^2)} \text{ for } j' \neq j, \text{ and } \check{U}_{j+1}^{(i)} = \frac{X_j^{(i)} X_{j\check{\bullet}}^{(i)}}{\check{E}(X_{j\check{\bullet}}^2)} - 1.$$

3. (Define the intercept estimator) The partial information (PI) estimator is

$$\{\hat{\beta}_{PI}\}_j = \bar{W}_j - \hat{a}\bar{U}_1 - \sum_{j'=1}^p \hat{b}_{j'} \bar{U}_{j'+1},$$

where $\hat{a}, \hat{b}_{j'}$ are the regular LSE of the regression model

$$W_j^{(i)} = \beta_j + a\check{U}_1^{(i)} + \sum_{j'=1}^p b_{j'} \check{U}_{j'+1}^{(i)} + \check{\delta}_j^{(i)}, \text{ for } i = 1, \dots, n.$$

The methodology is built from standard least squares procedures. An R code that implements the algorithm and also computes estimates of the variance as in Section 6.1 below, is available at the homepage of the first author.

4 Simplicity and semi-parametric efficiency

This section describes an optimality property of $\hat{\beta}_{TI}$ and discusses its relation to the estimator of Chakraborty and Cai (2018) their notion of semi-parametric efficiency.

4.1 Semi-parametric efficiency

To define a semi-parametric efficient estimator in this context, we rewrite model (15) as

$$Y = \alpha + \beta^T \mathbf{X} + \underbrace{\eta + \varepsilon}_{=\delta}, \quad (21)$$

where $\eta = \eta(\mathbf{X}) = E(Y|\mathbf{X}) - \alpha - \beta^T \mathbf{X}$ and $\varepsilon = Y - \alpha - \beta^T \mathbf{X} - \eta$ and assume that the marginal distribution of \mathbf{X} is known (total information). In this model, the semi-parametric efficiency bound for the variance is $\{E(\bar{\mathbf{X}}\bar{\mathbf{X}}^T)\}^{-1} E(\bar{\mathbf{X}}\bar{\mathbf{X}}^T \varepsilon^2) \{E(\bar{\mathbf{X}}\bar{\mathbf{X}}^T)\}^{-1}$ (see Chakraborty and Cai (2018) for more details).

Here is a general description of the approach of Chakraborty and Cai, which is based on the idea of imputation of the Y 's in the unlabeled sample. Specifically, when Y_i is missing, define $\hat{Y}_i = \hat{m}(\mathbf{X}_i)$ for $i = n + 1, \dots, n + m$, where \hat{m} is an estimate of $m(\mathbf{X}) = E(Y|\mathbf{X}) = \alpha + \beta^T \mathbf{X} + \eta(\mathbf{X})$ based on the labeled sample. They suggest to use the LSE of the unlabeled sample when using the imputed Y 's, i.e., based on $(\mathbf{X}_n, \hat{Y}_n), \dots, (\mathbf{X}_{n+m}, \hat{Y}_{n+m})$. When \hat{m} converges to m , the resulting estimator is semi-parametric efficient. However, in some situations, e.g., when $m(\mathbf{X})$ has a complicated parametric form and n is not very large, it is difficult to efficiently estimate it with the data set in hand. Below we review the pros and cons of this approach with respect to ours.

4.2 Simplicity

We now discuss the relation of our approach to semi-parametric efficiency and introduce an optimality property of our approach, which we call ‘‘simplicity’’. In order to avoid unnecessary technicalities, we consider here only the one-dimensional case under full information. Explicitly, let $(X, Y) \sim G$ where

$$\text{Var}_G(X) \in (0, \infty), \quad E_G(XY)^2 < \infty \quad \text{and} \quad E_G(X) = E_G(X^2) - 1 = 0; \quad (22)$$

the latter assumption means that the first and second moments of X are known. The sub-script G is used to emphasize that the expectation is taken over G as we will consider below several such G 's. In this context, the semi-parametric efficiency bound for the variance is $E_G(\varepsilon^2 X^2)$.

We argue that $\hat{\beta}_{TI}$ is the simplest (to be defined shortly) estimator within the framework of intercept model that dominates the LSE, and we also describe how to construct an estimator that achieves the semi-parametric efficiency bound, which is different from the estimator of Chakraborty and Cai. To this end, consider a basis of functions $\{q_k(X)\}_{k=1}^\infty$ that satisfies

$$\text{for all } K \in \mathbb{N} : \quad E_G[q_1(X)] = \dots = E_G[q_K(X)] = 0,$$

$$E_G \left[\{1, q_1(X), \dots, q_K(X)\}^T \{1, q_1(X), \dots, q_K(X)\} \right] \text{ is invertible,}$$

$$\text{and } \{XY, q_1(X), \dots, q_K(X)\} \text{ has finite second moments under } G. \quad (23)$$

The assumption that $E_G[q_1(X)] = \dots = E_G[q_K(X)] = 0$ is justified since the marginal distribution of X under G is assumed known.

Consider an intercept model based on $q_1(X), \dots, q_K(X)$

$$XY = \beta + \sum_{k=1}^K b_k q_k(X) + \tilde{\delta}_K. \quad (24)$$

When (23) holds, then b_1, \dots, b_K in (24) are well defined as the parameters of the best linear predictor. Since $E_G[q_1(X)] = \dots = E_G[q_K(X)] = 0$, the intercept is $E_G(XY) = \beta$.

Suppose we have n iid observations $(X_1, Y_1), \dots, (X_n, Y_n)$ from G . Let $\hat{\beta}_{LSE}$ be the least squares estimate based on n observations from Model (1) and let $\hat{\beta}_K$ be the intercept LSE estimate based on n observations from Model (24). By the ‘‘Sandwich’’ Theorem (4),

$$\sqrt{n} \left(\hat{\beta}_{LSE} - \beta \right) \xrightarrow{\mathcal{D}} N \left(0, E_G(\delta X)^2 \right) \text{ and } \sqrt{n} \left(\hat{\beta}_K - \beta \right) \xrightarrow{\mathcal{D}} N \left(0, E_G \tilde{\delta}_K^2 \right).$$

The following theorem states that $\hat{\beta}_K$ is better (smaller asymptotic variance) than $\hat{\beta}_{LSE}$ for every distribution G , iff $X \in \text{span}\{q_1(X), \dots, q_K(X)\}$ and $X^2 - 1 \in \text{span}\{q_1(X), \dots, q_K(X)\}$.

Theorem 3.

(Sufficiency) If $X \in \text{span}\{q_1(X), \dots, q_K(X)\}$ and $X^2 - 1 \in \text{span}\{q_1(X), \dots, q_K(X)\}$, then $\hat{\beta}_K$ is better than $\hat{\beta}_{LSE}$ for ‘‘every’’ G , i.e., $E_G \tilde{\delta}_K^2 \leq E_G(\delta X)^2$ for every distribution G that satisfies (22) and (23).

(Necessity) If $X \notin \text{span}\{q_1(X), \dots, q_K(X)\}$ or $X^2 - 1 \notin \text{span}\{q_1(X), \dots, q_K(X)\}$, then there exists distribution G that satisfies (22) and (23), and where $\hat{\beta}_{LSE}$ is better than $\hat{\beta}_K$, i.e., $E_G(\delta X)^2 < E_G \tilde{\delta}_K^2$.

Theorem 3 implies that $\hat{\beta}_{TI}$ is ‘‘simple’’ in the sense that it is derived from the minimal intercept model that is better than $\hat{\beta}_{LSE}$ for every G that satisfies (22) and (23). In this sense, $\hat{\beta}_{TI}$ is the simplest estimator that is derived from an intercept model.

This ‘‘simplicity’’ property allows us to improve over the LSE without explicitly modeling the non-linear part $\eta(\mathbf{X})$. This is one advantage over Chakraborty and Cai’s approach, where the missing Y ’s are imputed using an estimate of $\eta(\mathbf{X})$. We expect that if $\eta(\mathbf{X})$ has a complicated form and n is not very large, our estimate may be advantageous. On the other hand, a successful approximation of $\eta(\mathbf{X})$, may lead to an estimator of β with smaller variance, close to the semi-parametric bound. This is demonstrated and further discussed in Section 5.4 below.

The intercept model framework can also be used to construct an estimator that achieves the semi-parametric bound. To this end, assume for the moment that for some K

$$\alpha X + \beta X^2 + \eta X \in \text{span}\{q_1(X), \dots, q_K(X)\}. \tag{25}$$

Then Model (21) implies that $E_G(\tilde{\delta}_K^2) = E_G(\varepsilon^2 X^2)$ and therefore $\hat{\beta}_K$ is semi-parametric efficient. Generally, (25) is not likely to hold precisely, but if $\{q_k(X)\}_{k=1}^\infty$ is rich enough, then for large enough K , (25) holds true up to some small enough approximation error. In this case, one can increase K slowly with n and then achieve semi-parametric efficiency. This procedure is made precise in Theorem 7 of Zhang et al. (2017).

4.3 Other modifications

This addition of extra basis terms is similar to the typical practice of adding extra terms in the regression model, but it is better in a critical way: expanding the basis in (24) improves estimation of β uniformly

over G , whereas adding orthogonal terms to the regression model can sometimes harm it. To observe the latter fact, consider the one dimensional case (Model (12)) and the total information scenario where $E(X) = 0$ and $E(X^2) = 1$. One could modify the model to

$$Y = \alpha + \beta X + \gamma f(X) + \tilde{\delta}, \quad (26)$$

where $f(X)$ is some function of X that satisfies $E(f(X)) = 0$ and $E(Xf(X)) = 0$ (otherwise α and β might change). In this case, the asymptotic variance for the LSE of β in the original model (12) is $E(X^2\delta^2)$ and in the modified model (26) it is $E(X^2\tilde{\delta}^2)$. For a given $f(X)$, it is possible, under mild conditions, to find δ such that the modified estimator has larger asymptotic variance. To construct one, fix $\gamma = 1$ and pick an arbitrary residual $\tilde{\delta} = \tilde{\delta}_0$ such that $E(X^2f(X)\tilde{\delta}_0)$ is nonzero (assuming such $\tilde{\delta}_0$ exists). Then in the original model (12), consider the family $\delta = \delta_c = f(X) + c\tilde{\delta}_0$ for variable c . The difference of the asymptotic variances between the new LSE and the old LSE is

$$E(X^2c^2\tilde{\delta}_0^2) - E(X^2\delta_c^2) = -E(X^2f^2(X)) - 2cE(X^2f(X)\tilde{\delta}_0).$$

The variable c can be chosen to make this difference positive, yielding a model in which the ordinary least squares estimation has been harmed.

To sum up the findings of this section we conclude that the intercept model framework derives a large class of estimates that are better than $\hat{\beta}_{LSE}$. Our suggested estimator $\hat{\beta}_{TI}$ is minimal within this class and a semi-parametric efficient estimator can be achieved by considering a sequence of increasing models. Compared to Chakraborty and Cai's estimator it is especially useful when the non-linear part is difficult to model.

5 A simulation study

We compare the performance of the Partial Information (PI) and Total Information (TI) estimators against the Least Squares Estimate (LSE) across a wide range of settings.

5.1 Toy example

We start by studying the following toy model

$$Y = \alpha X^2 + \beta X + \varepsilon = \alpha + \beta X + \underbrace{\alpha(X^2 - 1) + \varepsilon}_{=\delta}, \quad (27)$$

where X and ε are i.i.d $N(0,1)$. Under model (27), the linear and non-linear part are determined by β and α separately and the linear coefficient β does not affect the residual δ . In this case, it is easy to calculate

the asymptotic variance of the estimates explicitly

$$\sigma_{LSE}^2 = 10\alpha^2 + 1, \quad \sigma_{TI}^2 = 6\alpha^2 + 1, \quad \sigma_{diff}^2 = 4\alpha^2.$$

Considering the excess risk, then by Proposition 1 the ratio of excess risks of $\hat{\beta}_{PI}$ ($\hat{\beta}_{TI}$, respectively) and $\hat{\beta}_{LSE}$ converges to $\frac{2\alpha^2+1+\sigma_{PI}^2}{2\alpha^2+1+\sigma_{LSE}^2}$ ($\frac{2\alpha^2+1+\sigma_{TI}^2}{2\alpha^2+1+\sigma_{LSE}^2}$, respectively). The limits equal 1 in the linear case ($\alpha = 0$) and approach 0 when the non-linear part is dominant ($\alpha \rightarrow \infty$).

Figure 1 summarizes 10,000 simulations of model (27) with $\beta = 1$ and different values of α , n ; for the PI case we used $m = 2n$. For all the scenarios we found that $\hat{\beta}_{TI}$, $\hat{\beta}_{PI}$ have smaller excess risk when n is large enough. As α increases, the departure from linearity is more significant and the ratio of the excess risks is smaller. For small n , LSE is superior for all scenarios as it is a simpler estimate. When the model is close to linear ($\alpha = 1/4$), the new estimates are better for $n \geq 200$ and for the other values of α the new estimates are better for $n \geq 70$. For small values of α the limiting excess risk is close to the actual risk even when $n \approx 300$ but this does not hold true for larger values of α . In short, we found that the new estimates are better for large n , and are much better when the non-linear part is significant.

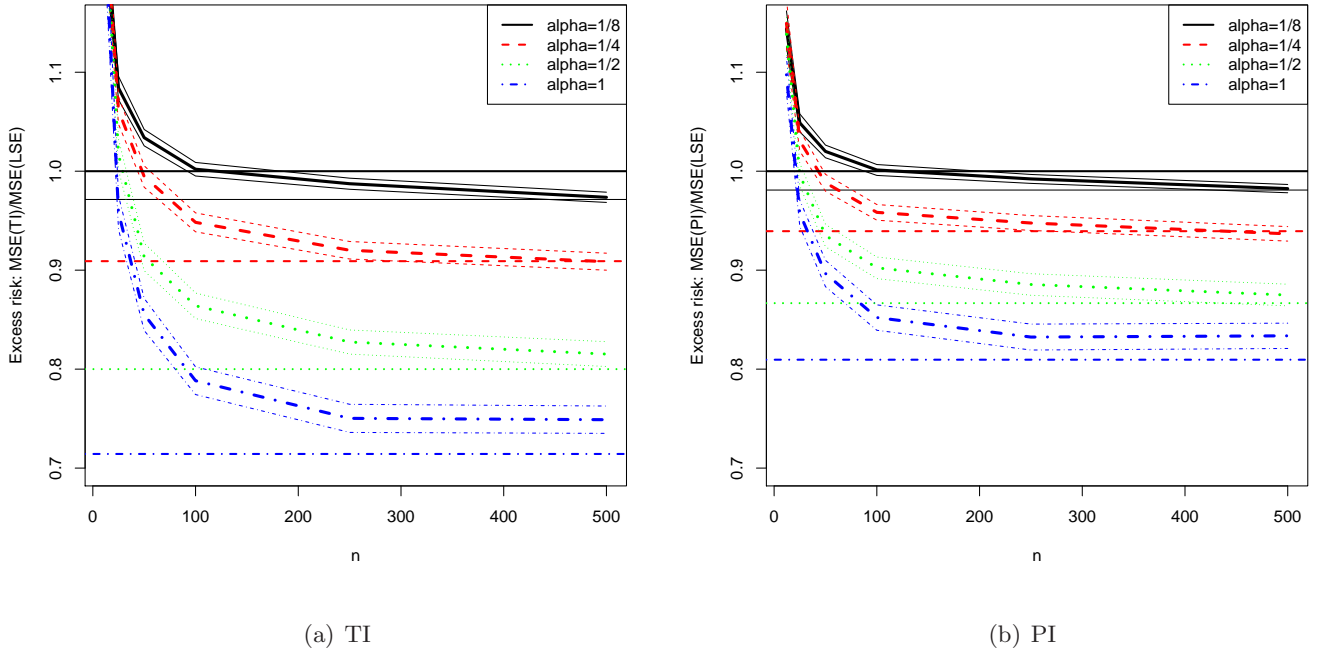


Figure 1: Estimates of the excess risk ratio for the TI and PI estimates based on 10,000 simulations of model (27) with $\beta = 1$ and different values of α , n ; for the PI case we used $m = 2n$. A confidence interval based on two standard deviations is also plotted. The horizontal lines represents the limiting excess risk.

5.2 Simulation Parameters

We now compare our new estimates to LSE for a broader variety of scenarios. Estimates of the Excess Risk Ratio (ERR) for PI and TI were obtained for every combination of the following parameter choices:

n (size of labeled dataset)	12, 25, 50, 100, 250, 500
m (size of unlabeled dataset)	$n, 2n$
p (number of predictors)	1, 4
Distribution of X	Gaussian, Lognormal, Exponential, Cubed Gaussian
Errors	Gaussian, $N(0, e^{2\ X\ })$
$E[Y X]$	$X, e^X, X^3, \sqrt{ X }$

For $p = 4$, $E[Y|X]$ was the sum over j of the chosen function of each X_j . Fixing each of the 768 parameters settings, many sample datasets were generated, along with a large test dataset ($n = 100,000$). For each sample dataset, PI, TI, and LSE were fit, (calculating TI using the PI method with $m \geq 500n$.) Next, an estimate of expected MSE was obtained using the test dataset. Along with the best linear fit for the test dataset (given by least squares regression), these determine the Excess Risk Ratio for PI (ERR_{PI}) as follows:

$$ERR_{PI} = \frac{MSE_{PI} - MSE_{BLF}}{MSE_{LSE} - MSE_{BLF}},$$

where

MSE_{PI} is the expected MSE of the PI estimator (estimated from mean performance on the test dataset), MSE_{LSE} is the expected MSE of the LSE estimator, and MSE_{BLF} is the MSE of the best linear fit of the test dataset.

An analogous calculation was performed to calculate ERR_{TI} . Additional datasets were sampled to improve estimates of MSE_{PI} and MSE_{LSE} , until standard errors for ERR_{PI} and ERR_{TI} fell below 1% (as estimated by the delta method), or a maximum of 100,000 sample datasets was generated.

5.3 Simulation Results

Table 3 provides the proportions of scenarios where the PI estimate yields a statistically significant smaller excess risk than the LSE and the proportion where the opposite holds true, for different p and n . Statistical significance is measured by two standard errors. The results demonstrate that the PI estimate outperforms the LSE across a wide range of scenarios for large n , and the proportion increases with n . When $n = 500$, LSE is significantly better only for about 10% of the scenarios we studied. More detailed comments follow:

Table 3: Proportions of scenarios where ERR_{PI} is significantly smaller than 1 and where it is significantly larger than 1.

	p \ n	n					
		12	25	50	100	250	500
$ERR_{PI} < 1$ significantly	1	0.297	0.344	0.469	0.594	0.656	0.609
	4	0.047	0.234	0.391	0.438	0.594	0.547
$ERR_{PI} > 1$ significantly	1	0.344	0.234	0.172	0.125	0.125	0.109
	4	0.578	0.391	0.234	0.188	0.188	0.125

1. A proposed rule of thumb when $p = 1$: As indicated by Table 3 across a wide range of parameter settings, PI performs almost as well as LSE when $n \geq 100$, $m \geq 100$ and linear $E[Y|X]$, and better than the LSE for nonlinear $E[Y|X]$ (that is, $ERR < 1$). Furthermore, PI's underperformance is mild when $E[Y|X]$ is linear. Thus, assuming that the scenarios we studied are representative of reality, we recommend using the PI method when the following conditions are satisfied, as a rule of thumb:

- (a) $n \geq 100$
- (b) $m \geq 100$
- (c) $E[Y|X]$ could be non-linear

2. If $p = 1$, $n > 100$, $m > 100$, and $E[Y|X]$ is one of the non-linear functions tested, then both PI and TI appear to outperform LSE; furthermore, the margin of outperformance widens as n increases (See Figure 2 (a), and Figure 2 (b), respectively). In particular, PI does much better for $E[Y|X] = X^3$ and e^X , as compared to \sqrt{X} . Intuitively, this results from the larger non-linear moments of X^3 and e^X , causing Σ_{diff} to be large (as defined in Theorem 2, part ii)

3. Increasing m , the size of the unlabeled dataset, improves performance of PI. Indeed, comparing PI where $m = 2n$ in Graph X against TI ($m \geq 500n$) in Graph Y, one observes that greatly increasing the pool of unlabeled samples can reduce Excess Risk, in some cases on the order of 5 – 10 per cent (Compare Figures 2 (a) and (b)). But, a word of caution: if n is small, then LSE may still outperform both PI and TI, even when m is large and $E[Y|X]$ is non-linear.

4. If n is small, or if $E[Y|X]$ is linear, then LSE performs better than TI and PI. However, $ERR \rightarrow 1$ as $n \rightarrow \infty$, holding other parameters fixed (See Figures 2 (a) and (b)). This result is consistent with the asymptotic agreement of PI, TI, and LSE when $E[Y|X]$ is linear (See Theorem 2)






Figure 2: Estimates of the excess risk ratio for the TI and PI estimates, for several possible functions of $E[Y|X]$. In PI estimates (a), $m = 2n$. In TI estimates (b), the PI method is used with $m \geq 500n$. In all cases, X is Gaussian, errors are Gaussian, and $p = 1$. Confidence intervals are ± 2 standard errors.


5. Extra heteroskedastic noise appears to decrease the relative advantage of PI. See Figure 3 (a).
6. Somewhat similarly, increasing the number of parameters to 4 (with $E[Y|X]$ being the sum over j of the given function of X_j) also decreases the relative advantage of PI. See Figure 3 (b).
7. If X is not Gaussian, PI still does well, especially for large n . See Figure 3 (c) for an example where X is exponential. (Note that standard errors may be understated, especially in the case $E[Y|X] = \exp(X)$, partly due to the large moments of Y .)



(a) Heteroskedastic Errors



(b) $p = 4$



(c) X Exponential

Figure 3: Estimates of the excess risk ratio of PI estimates, across several possible functions of $E[Y|X]$, with tweaks to the previous settings of Figure 2 (a): these were $m = 2n$, X Gaussian, $p = 1$, and Gaussian errors. The following is changed in each plot: Figure (a): Errors are heteroscedastic, distributed $N(0, e^{2|x|})$. Figure (b): $p = 4$, where $E[Y|X]$ is the sum of the given relationship in each X -variable. Figure (c): X is distributed exponentially. Confidence intervals are ± 2 standard errors.

5.4 Comparison to Chakraborty and Cai

In this section we describe a simulation study that compares $\hat{\beta}_{TI}$ with the estimator of Chakraborty and Cai (2018), using R code they generously provided. They suggest two candidate smoothing methods, but in both their simulation results and ours, one method (kernel machine regression, or in their notation, $T := \text{KM}$) performs better than the other ($T := \text{KS}_{2,P_2}$). Therefore we report the results only of the former. We denote it “EASE” (Efficient and Adaptive Semi-Supervised Estimator). The kernel machine regression procedure is computationally intensive, especially when n is not small, and therefore we did not compute the EASE estimator for all scenarios considered in Section 5.2, but instead chose several interesting ones.

In order to compare the two methods we ran a simulation study, where the dimension is $p = 4$, the number of labeled observations is $n = 100$, and the unlabeled sample is much larger $m = 10,000$ (Chakraborty and Cai’s estimator is constructed for cases where $m \gg n$). The simulations were repeated 1000 times. We considered four scenarios for the joint distribution of (\mathbf{X}, Y) ; in all scenarios the distribution of \mathbf{X} and ε is iid $N(0, 1)$.

- Scenario 1: $Y = (1, 1, 1, 1)\mathbf{X} + (1, 1, 1, 1)\mathbf{X}^2 + \varepsilon$ (where $\mathbf{X}^2 = (X_1^2, X_2^2, X_3^2, X_4^2)^T$ and similar notation is used below for other functions).
- Scenario 2: $Y = (1, 1, 1, 1)\mathbf{X} + (0.3, 0.3, 0.3, 0.3)\mathbf{X}^2 + \varepsilon$.
- Scenario 3: $Y = (1, 1, 1, 1)\mathbf{X} + (1, 1, 1, 1)(\mathbf{X}^3 - \mathbf{X}^2 + \exp(\mathbf{X})) + \varepsilon$.
- Scenario 4: $Y = (1, 1, 1, 1)\mathbf{X} + (0.3, 0.3, 0.3, 0.3)(\mathbf{X}^3 - \mathbf{X}^2 + \exp(\mathbf{X})) + \varepsilon$.

In Scenarios 1 and 2 the non-linear part has a squared form; they differ in the level of non-linearity. On the other hand, Scenarios 3 and 4 have a complex non-linear part, which is difficult to estimate given the relatively small number of observations.


Table 4: Mean (std) of the MSE of $\hat{\beta}_{LSE}$, $\hat{\beta}_{PI}$, and EASE.

	$\hat{\beta}_{LSE}$	$\hat{\beta}_{PI}$	EASE
Scenario 1	0.174 (0.004)	0.137 (0.003)	0.081 (0.003)
Scenario 2	0.025 (0.001)	0.023 (0.001)	0.022 (0.001)
Scenario 3	0.985 (0.031)	0.773 (0.026)	0.961 (0.028)
Scenario 4	0.106 (0.005)	0.076 (0.002)	0.097 (0.003)

To match the convention of Chakraborty and Cai (2018), in this section we report results in terms of MSE of β rather than excess risk. (The ERR defined in 5.2 can be recovered asymptotically as the ratio of the MSE of β to the MSE of β_{LSE} , due to Proposition 1). The mean MSE and standard errors of the simulations are given in Table 4. EASE outperformed all other methods in Scenario 1; its MSE is about half than the LSE. In scenario 2 all methods are more or less the same since the non-linear part is relatively small. Scenarios 3 and 4 suit better our framework since $\eta(\mathbf{X})$ is a complex function and is difficult to estimate. Indeed, in these scenarios, $\hat{\beta}_{PI}$ improves over the LSE and also outperforms EASE. Notice that the relative improvement of $\hat{\beta}_{PI}$ remains the same across Scenarios 3 and 4 (the MSE in Scenario 4 is about 10% of the MSE of Scenario 3 in all estimates).

We also report two typical examples of how the relative performance of the estimators can depend on n . We considered two scenarios for the joint distribution of (\mathbf{X}, Y) , which are chosen from those of Section 5.2:

- Scenario 5: $\mathbf{X} \sim N(0, 1)$ i.i.d. and $Y = (1, 1, 1, 1)\mathbf{X} + (.3, .3, .3, .3)\mathbf{X}^3 + \varepsilon$
- Scenario 6: $\mathbf{X} \sim Exp(1)$ i.i.d. and $Y = (1, 1, 1, 1)\mathbf{X}^{1/2} + \varepsilon$



(a) Scenario 5

(b) Scenario 6

Figure 4: Confidence intervals are ± 2 simulation standard errors

In both scenarios, ε is iid $N(0, 1)$ and the dimension is $p = 4$. For labeled dataset sizes $n = 12, 25, 100, 250, 500$, and unlabeled dataset sizes $m = 20n$, for both $\hat{\beta}_{PI}$ and $\hat{\beta}_{EASE}$ we computed via

simulation the MSE's of $\hat{\beta}_{PI}$ and $\hat{\beta}_{EASE}$ and standardized by the simulated MSE of $\hat{\beta}_{LSE}$. The ratios are displayed in Figure 4, with upper and lower intervals of ± 2 simulation standard errors based on the delta method. Both estimators can be worse than the LSE for small n . Both scenarios show an interesting pattern: PI's MSE compared to LSE improves uniformly with n , whereas for EASE the ratio kinks upward at small n and then decreases. In Scenario 5, PI does significantly better than both EASE and LSE for moderate n , but for large n EASE eventually catches up. In Scenario 6, the nonlinearity is mild and both PI and EASE struggle; we do not observe a range where PI is simultaneously better than both EASE and the LSE, and PI and EASE only beat the LSE for large n .

In practice, one does not know the complexity or severity of the nonlinearity relative to n , and thus it is unclear which estimator is best. Therefore, it might be advisable to try all three methods. Since each method estimates its standard error, one can try to evaluate the proper estimator for the specific data set in hand.

6 Estimating the standard errors of the estimates

In this section we study estimation of the asymptotic covariance matrix. We suggest three estimates for β_{PI} in Section 6.1 and study their performance by simulation (Section 6.2). Section 6.3 deals with the estimation of the prediction error (excess risk).

6.1 Three estimates of Σ_{PI}

We proved that under certain conditions,

$$\sqrt{n}(\hat{\beta}_{LSE} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{LSE}), \sqrt{n}(\hat{\beta}_{PI} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{PI}) \text{ and, } \sqrt{n}(\hat{\beta}_{TI} - \beta) \xrightarrow{\mathcal{D}} N(0, \Sigma_{TI}).$$

In these expressions the (asymptotic) variability of the estimates are determined by the matrices Σ_{LSE} , Σ_{PI} and Σ_{TI} . These matrices are unknown and therefore need to be estimated. In this section we consider only estimation of Σ_{LSE} and Σ_{PI} . Estimation of the asymptotic variance yields prediction error estimates or L_2 risk in a standard fashion. In section 6.3 we discuss estimation of the excess risk.

For Σ_{LSE} we follow Buja et al. (2017) who show that the standard estimates are inconsistent when non-linearity is present and suggest two alternative estimates for the variance. The parametric sandwich estimator for the asymptotic variance of $\hat{\beta}_{LSE}$ is

$$\hat{AV}_{param}(\hat{\beta}_{LSE}) = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \hat{D} \mathbb{X} (\mathbb{X}^T \mathbb{X})^{-1},$$

where \mathbb{X} is the design matrix and \hat{D} is a diagonal matrix with $\hat{D}_i = \hat{\delta}_i^2$ the standard residual estimator. The second estimate is derived from a pairs bootstrap where a pair (\mathbf{X}^*, Y^*) is sampled N_{BS} times from

the empirical joint distribution of (\mathbf{X}, Y) , yielding a sample (under the empirical measure) of size N_{BS} of $\hat{\beta}_{LSE}$. The resulting estimate of the variance is denoted by $\hat{AV}_{BS}(\hat{\beta}_{LSE})$.

For estimating the variance of Σ_{PI} , we consider three estimates described as follows:

1. **Parametric:** The asymptotic variances of $\hat{\beta}_{LSE}$ and $\hat{\beta}_{PI}$ are

$$AV(\hat{\beta}_{LSE}) = Cov(\tilde{\delta}) + Cov(\mathbf{V}), \quad AV(\hat{\beta}_{PI}) = Cov(\tilde{\delta}) + \nu Cov(\mathbf{V}). \quad (28)$$

To estimate the variance of $Cov(\tilde{\delta})$ notice that for $j = 1, \dots, p$ and $i = 1, \dots, n$,

$$\tilde{\delta}_j^{(i)} = \check{W}^{(i)} - \{\hat{\beta}_{PI}\}_j - \hat{a}\check{U}_1^{(i)} - \sum_{j'=1}^p \hat{b}_{j'}\check{U}_{j'+1}^{(i)} = \check{\delta}_j^{(i)} + o_p(1/\sqrt{n}),$$

and hence, $\overline{Cov}(\check{\delta}) \xrightarrow{p} Cov(\tilde{\delta})$, where \overline{Cov} is the empirical covariance based on the labeled n sample. Therefore, using (28) we obtain that a consistent estimator to $AV(\hat{\beta}_{PI})$ is

$$\left\{ \hat{AV}_{param}(\hat{\beta}_{PI}) \right\}_{j,j'} = \overline{Cov}(\check{\delta}_j, \check{\delta}_{j'}) + \nu \left(\left\{ \hat{AV}_{param}(\{\hat{\beta}_{LSE}\}) \right\}_{j,j'} - \overline{Cov}(\check{\delta}_j, \check{\delta}_{j'}) \right),$$

for $j \neq j'$ and for $j = 1, \dots, p$,

$$\left\{ \hat{AV}_{param}(\hat{\beta}_{PI}) \right\}_{j,j} = \frac{1}{n} \sum_{i=1}^n \left\{ \check{\delta}_j^{(i)} \right\}^2 + \nu \max \left(\left\{ \hat{AV}_{param}(\{\hat{\beta}_{LSE}\}) \right\}_{j,j} - \frac{1}{n} \sum_{i=1}^n \left\{ \check{\delta}_j^{(i)} \right\}^2, 0 \right).$$

2. **Bootstrap:** Let $\left\{ (\mathbf{X}_*^{(i)}, Y_*^{(i)}) \right\}_{i=1}^n$ be a sample from the empirical distribution of (\mathbf{X}, Y) (based on the labeled observations) and $\left\{ \mathbf{X}_*^{(i)} \right\}_{i=n+1}^{n+m}$ be a sample from the empirical distribution of \mathbf{X} (based on the unlabeled observations). Thus, $\hat{\beta}_{PI}^* = \hat{\beta}_{PI} \left(\left\{ (\mathbf{X}_*^{(i)}, Y_*^{(i)}) \right\}_{i=1}^n, \left\{ \mathbf{X}_*^{(i)} \right\}_{i=n+1}^{n+m} \right)$ is a sample under the empirical measure of $\hat{\beta}_{PI}$, which can yield an estimate of the variance of the estimator, denoted by $\hat{AV}_{BS}(\hat{\beta}_{PI})$.

3. **Variance bootstrap:** In the proof of Theorem 2 we showed that for $j = 1, \dots, p$,

$$\{\hat{\beta}_{LSE}\}_j = \beta_j + \frac{\frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)} \delta^{(i)}}{E(X_{j\bullet}^2)} + o_p(1/\sqrt{n}) \quad \text{and} \quad \{\hat{\beta}_{PI}\}_j = \beta_j + \frac{1}{n} \sum_{i=1}^n \tilde{\delta}^{(i)} + \frac{1}{n+m} \sum_{i=1}^{n+m} V_j^{(i)} + o_p(1/\sqrt{n}).$$

Hence, the definition of the vector \mathbf{V} implies that

$$\begin{aligned} \hat{\beta}_{LSE} - \hat{\beta}_{PI} &= \frac{1}{n} \sum_{i=1}^n -\mathbf{V}^{(i)} + \frac{1}{n+m} \sum_{i=1}^{n+m} \mathbf{V}^{(i)} + o_p(1/\sqrt{n}) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{V}^{(i)} \left(\frac{n}{n+m} - 1 \right) + \frac{1}{n+m} \sum_{i=n+1}^{n+m} \mathbf{V}^{(i)} + o_p(1/\sqrt{n}). \end{aligned}$$

Therefore, dropping the smaller order terms we obtain

$$\begin{aligned} nCov(\hat{\beta}_{LSE} - \hat{\beta}_{PI}) &\approx Cov(\mathbf{V}) \left(1 - \frac{n}{n+m} \right)^2 + \frac{nm}{(n+m)^2} Cov(\mathbf{V}) \\ &= Cov(\mathbf{V}) \left(1 - \frac{n}{n+m} \right) \rightarrow Cov(\mathbf{V}) (1 - \nu) = \Sigma_{LSE} - \Sigma_{PI}. \end{aligned}$$

Thus, a consistent estimator of the difference $\Sigma_{LSE} - \Sigma_{PI} =: \Delta$ can be obtained by bootstrapping $nCov(\hat{\beta}_{PI}^* - \hat{\beta}_{LSE}^*)$. We denote this estimator by $\hat{\Delta}$. A consistent estimate to $AV(\hat{\beta}_{PI})$ is

$$\hat{AV}_{VBS}(\hat{\beta}_{PI}) = \hat{AV}_{BS}(\hat{\beta}_{LSE}) - \hat{\Delta}.$$

This estimator has the advantage over the previous suggestions that it is always smaller than the estimated variance of the LSE (in the sense that the difference is positive definite), in a similar fashion to the asymptotic distribution.

In the next section we compare the different estimates through a simulation study.

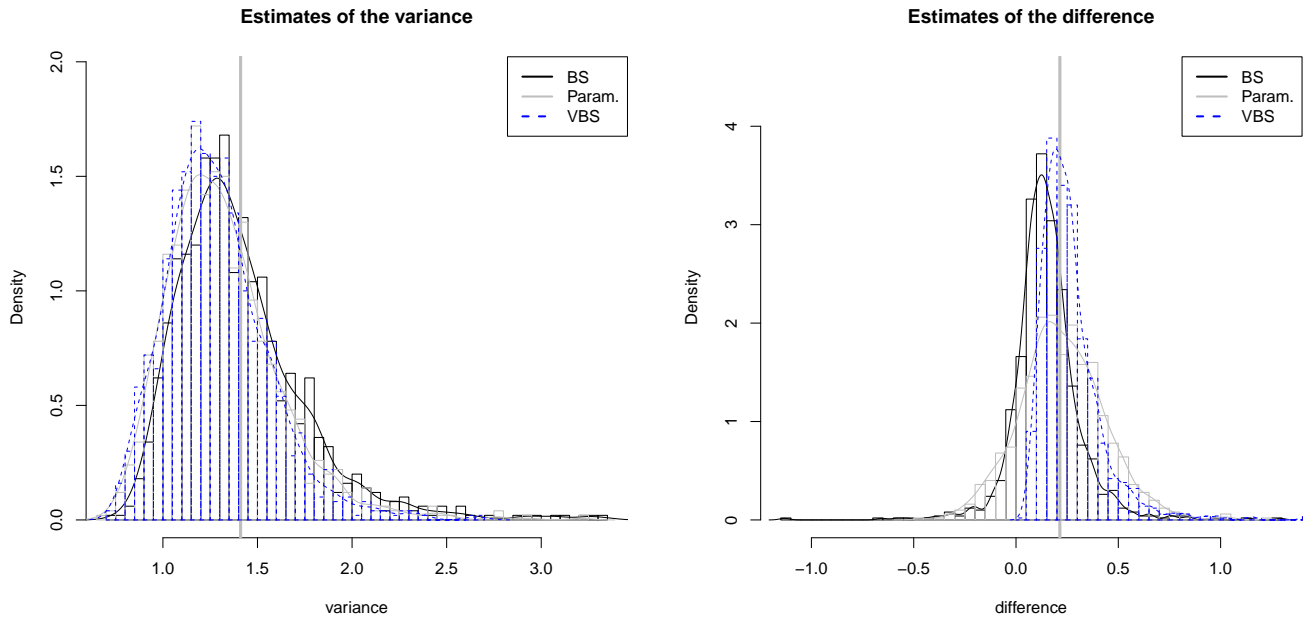
6.2 Simulations for the estimates of the variance

We now investigate the estimates of the variance under the toy model (27). We study the performance of the estimates for $\alpha = 1, 1/2, 1/4, 1/8$, where small α corresponds to little non-linearity. We considered $n = 250$ and $m = 1500$ similar to the numbers of Section 7. For each α we repeated the simulation 1000 times and also $N_{BS} = 1000$.

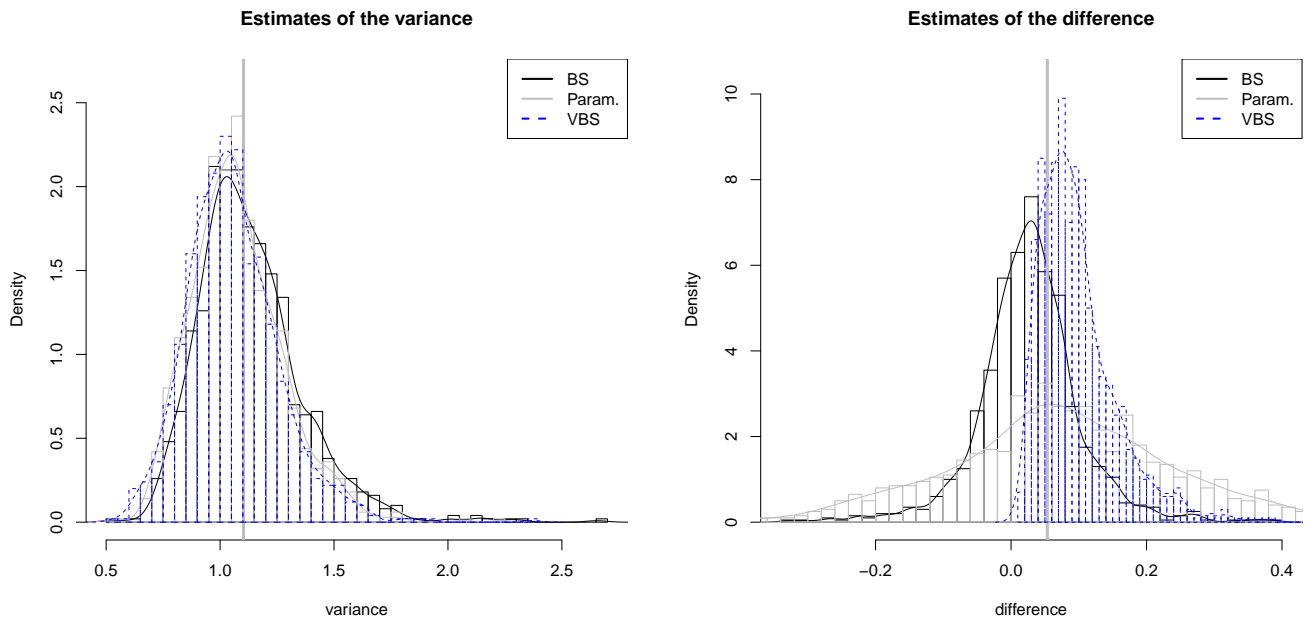
The simulations are summarized in Table 5 and in Figure 5. We find that the bootstrap estimate (henceforth BS) is more variable than the parametric (henceforth PARM) and the variance bootstrap (henceforth VBS). PARM and VBS are comparable but for small α 's for estimating the difference, the histogram of VBS is narrower around the true value. When $\alpha = 1/8$, the estimated difference under BS (PARM) is negative 32.9% (34.7%) of the simulations, while for VBS it is never negative. However, since VBS cannot be negative, the mean of the simulations is upwards biased. In short, we find that PARM and VBS outperform BS and VBS has the advantage of never being smaller than $\hat{AV}(\hat{\beta}_{LSE})$.

Table 5: The mean (std) of the estimates of the asymptotic variance $AV(\hat{\beta}_{PI})$ and the difference $AV(\hat{\beta}_{LSE}) - AV(\hat{\beta}_{PI})$.

α	Variance				Difference			
	Bootstrap	Parametric	Variance BS	true	Bootstrap	Parametric	Variance BS	true
1	7.65 (3.71)	6.86 (3.11)	6.58 (2.40)	7	2.63 (2.53)	3.51 (1.53)	3.70 (2.13)	3.43
1/2	2.65 (0.97)	2.42 (0.82)	2.36 (0.70)	2.5	0.64 (0.53)	0.88 (0.45)	0.93 (0.46)	0.86
1/4	1.43 (0.40)	1.33 (0.33)	1.31 (0.31)	1.37	0.15 (0.18)	0.24 (0.22)	0.27 (0.16)	0.21
1/8	1.11 (0.23)	1.06 (0.20)	1.03 (0.19)	1.09	0.02 (0.07)	0.06 (0.17)	0.10 (0.06)	0.05



(a) $\alpha = 1/4$



(b) $\alpha = 1/8$

Figure 5: Histograms of the estimates (Bootstrap, Parametric, Variance BS) of the asymptotic variance $AV(\hat{\beta}_{PI})$ and the difference $AV(\hat{\beta}_{LSE}) - AV(\hat{\beta}_{PI})$. The true asymptotic value is illustrated by the vertical gray line.

6.3 Estimating the prediction error

The three estimates of Σ_{PI} can be used to estimate the prediction error in a standard fashion. Here, we describe how the variance bootstrap estimate together with the result of Proposition 1 provide an estimate of the excess risk.

By Proposition 1, the excess risk of $\hat{\beta}_{LSE}$ can be approximated by

$$E\left[Y - E(Y) - \sum_{j=1}^p \beta_j \{X_j - E(X_j)\}\right]^2 + \text{Trace}(\mathbf{M}\Sigma_{LSE}) = \text{Var}(Y) + \beta^T \mathbf{M} \beta - 2 \sum_{j=1}^p \beta_j \text{Cov}(Y, X_j) + \text{Trace}(\mathbf{M}\Sigma_{LSE})$$

and hence can be consistently estimated by

$$\hat{\sigma}_Y^2 + \hat{\beta}_{LSE}^T \hat{\mathbf{M}} \hat{\beta}_{LSE} - 2 \sum_{j=1}^p \{\hat{\beta}_{LSE}\}_j \overline{\text{Cov}}(Y, X_j) + \text{Trace}\{\hat{\mathbf{M}} \hat{A} V_{BS}(\Sigma_{LSE})\},$$

where $\hat{\sigma}_Y^2 = \frac{1}{n} \sum_{i=1}^n (Y^{(i)} - \bar{Y})^2$, $\hat{\mathbf{M}} = \frac{1}{n+m} \sum_{i=1}^{n+m} (\mathbf{X}^{(i)} - \bar{\mathbf{X}}) \{\mathbf{X}^{(i)} - \bar{\mathbf{X}}\}^T$ and $\overline{\text{Cov}}(Y, X_j) = \frac{1}{n} \sum_{i=1}^n (Y^{(i)} - \bar{Y})(X_j^{(i)} - \bar{X}_j)$ is the empirical covariance. Therefore, a consistent estimate to the excess risk ratio is

$$\widehat{ERR}_{PI} = \frac{\hat{\sigma}_Y^2 + \hat{\beta}_{LSE}^T \hat{\mathbf{M}} \hat{\beta}_{LSE} - 2 \sum_{j=1}^p \{\hat{\beta}_{LSE}\}_j \overline{\text{Cov}}(Y, X_j) + \text{Trace}\{\hat{\mathbf{M}}(\hat{A} V_{BS}(\Sigma_{LSE}) - \hat{\Delta})\}}{\hat{\sigma}_Y^2 + \hat{\beta}_{LSE}^T \hat{\mathbf{M}} \hat{\beta}_{LSE} - 2 \sum_{j=1}^p \{\hat{\beta}_{LSE}\}_j \overline{\text{Cov}}(Y, X_j) + \text{Trace}\{\hat{\mathbf{M}} \hat{A} V_{BS}(\Sigma_{LSE})\}}, \quad (29)$$

where $\hat{\Delta}$ is defined by the Variance bootstrap.

7 Inferring homeless population

We now consider the Los Angeles homeless data set (Kriegler and Berk, 2010). Our purpose here is not to carefully analyze this data set but rather to demonstrate our new method and to compare it to the standard LSE. For our analysis, we used the same covariates as in Kriegler and Berk (2010); see this reference for more details about the data.

There are about 2000 census tracts in the Los Angeles county, and a sample was conducted in order to infer the homeless population. The sample consisted of two parts. First, tracts believed to have large numbers of homeless people were pre-selected and visited; there are about 240 such tracts. Second, a sample of about 260 tracts was chosen from the remainder by a random sampling technique and the homeless population was counted, leaving about 1500 tracts to be imputed.

Leaving aside the preselected tracts this is almost exactly our semi-supervised setting among the remaining population. We have $n = 261$ and $m = 1536$. The major difference from the setting in our introduction here is that sampling from the remaining population is without replacement and hence is not an i.i.d random sample. Since we have not discussed this type of sampling in detail, we will ignore this issue in the subsequent discussion which is for illustrative purposes. (The difference in sampling schemes does not affect the values of the various estimates, but does impact their standard errors.)

We compare the LSE of the random sample and the new estimate (PI) where the supervised part is the random sample (of size 261) and the unsupervised part consists of the 1536 tracts to be imputed. The resulting estimates, as well as their standard errors (SE) are given in Table 6. The SE of the LSE was computed using pairs bootstrap ($N_{BS} = 10^5$), and the SE of PI was computed using the above variance bootstrap method (the SE of the intercept of PI was estimated by pairs bootstrap).

We find that the SE's of PI are smaller by 15%-40% than LSE in the predictors MedianInc, PercVacant, PercCommercial, PercIndustrial. These predictors also demonstrated a discrepancy between the SE from linear models to bootstrap SE as reported in Buja et al. (2017) Section 2, indicating non-linearity in these predictors (a discrepancy was observed also in the predictor PercResidential, but the improvement in this predictor is relatively small). In summary, we found that the additional information on the distribution of the predictors in the unsupervised data provides estimates that are more accurate (smaller variance).

Consider the prediction problem of estimating the homeless population for the 1500 tracts. An estimate of the excess risk ratio, using (29), is $\widehat{ERR}_{PI} = 0.904$, i.e., improvement of about 10%. An estimate of the differences of prediction errors as in (10) and (11) is $\text{Trace}(\hat{\Delta}\hat{\mathbf{M}})/n = 15.2$, while an estimate of the prediction error of the mean is

$$\frac{1}{m} \sum_{i=1}^m \vec{\mathbf{X}}_{n+i}^T \hat{\Sigma}_{LSE} \vec{\mathbf{X}}_{n+i} / n = 42.8.$$

Since $15.2/42.8 = 0.355$ the improvement in prediction risk is about 36%. When considering the prediction of Y itself as in (11) then the improvement is much smaller, about 1.5%. The reason for this difference is that the predictors we considered are all quite weak, and hence most of prediction error comes from the variability in the distribution of Y given \mathbf{X} rather than the variability in estimating the β 's.

Table 6: The estimates $\hat{\beta}_{LSE}$, $\hat{\beta}_{PI}$ and their standard errors.

	$\hat{\beta}_{LSE}$	\hat{SE}_{LSE}	$\hat{\beta}_{PI}$	\hat{SE}_{PI}	$\frac{SE_{LSE}}{SE_{PI}}$
Intercept	13.138	11.822	14.758	11.184	0.927
MedianInc (\$K)	-0.065	0.056	-0.080	0.044	1.284
PercVacant	1.449	0.707	1.583	0.514	1.374
PercMinority	0.060	0.105	0.063	0.099	1.058
PercResidential	-0.070	0.100	-0.078	0.094	1.072
PercCommercial	0.446	0.354	0.335	0.300	1.179
PercIndustrial	0.101	0.188	0.202	0.143	1.317

8 Discussion

In this work we showed that β can be better estimated when additional information on the predictors \mathbf{X} is available. The key idea is to replace the regression model

$$Y = \alpha + \beta_1 X_1 + \cdots + \beta_p X_p + \delta, \quad (30)$$

with p regressions of the form

$$W_j = \beta_j + aU_1 + \sum_{j'=1}^p b_{j'} U_{j'+1} + \tilde{\delta}_j, \quad (31)$$

for $j = 1, \dots, p$, where W_j, U_1, \dots, U_{p+1} is defined in (18). Regression (31) can be used only when the first and second moments of \mathbf{X} are known. We showed that the intercept estimator of (31) is a better estimate of β_j , in terms of smaller asymptotic variance, than the standard LSE of (30). Furthermore, we showed that even if the second moments of \mathbf{X} are not known exactly then improvement can be made over the standard LSE provided that the unlabeled sample is non-negligible with respect to the labeled sample, i.e., m/n is bounded away from 0.

Here we considered the classical framework where p is fixed and $n \rightarrow \infty$. Zhang et al. (2017) study the semi-supervised mean estimation problem where p, n both go to infinity but p grows relatively slowly, i.e., $p = O(\sqrt{n})$. The findings of Zhang et al. (2017) indicate that our results can be extended to the latter case. However, our method obviously requires that $p \ll n$ and for high-dimensional regression models where $p > n$, a different approach is needed. High-dimensional misspecified regression models are the topic of a recent manuscript by Bühlmann and van de Geer (2015).

In this work we did not study variable selection and considered the covariates as given. However, improved asymptotic performance may be achieved by including functions of the existing covariates as additional new covariates. Section 3 of Zhang et al. (2017) discusses how this can be done for the problem of semi-supervised estimation of a mean. On the other hand, when the labeled sample, n , is small, then it may be better not to include all covariates. We hope to study the problem of variable selection in this context in a future research.

A possible extension of these results is to generalized linear models, where the label variable (Y) takes a small number of values and the regression model reduces to a classification problem. Our hope is that our method can be extended to these cases and improvement can be made over naive classifiers that consider only the labeled data.

In summary, in this work we considered the framework where the best linear predictor is of interest even if $E(Y|\mathbf{X})$ is not linear. Under this framework, additional information on the distribution of \mathbf{X} is useful to construct better estimates than the standard estimates. We believe that this framework is of practical importance and also can lead to interesting statistical issues, which are yet to be studied.

9 Proofs

Lemma 1. *We have that*

$$EL(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) = E\left[\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{\bar{X}_j - E(X_j)\} - E(Y)\right]^2 + E\left\{(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{M}(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})\right\}. \quad (32)$$

Proof.

First,

$$\begin{aligned} R(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) &= E\left(Y^* - \tilde{\alpha} - \tilde{\boldsymbol{\beta}}^T \mathbf{X}^*\right)^2 - E\left(Y^* - \alpha - \boldsymbol{\beta}^T \mathbf{X}^*\right)^2 \\ &= E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\} \left\{2Y^* - (\alpha + \tilde{\alpha}) - (\boldsymbol{\beta} + \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\} \\ &= E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\}^2 + 2E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\} \left\{Y^* - \alpha - \boldsymbol{\beta}^T \mathbf{X}^*\right\} \\ &= E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\}^2, \end{aligned}$$

where the last equality holds true since $Y^* - \alpha - \boldsymbol{\beta}^T \mathbf{X}^*$ is orthogonal to $X^*, 1$. We can further write

$$\begin{aligned} E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{X}^*\right\}^2 &= E\left[(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T E(\mathbf{X}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \{\mathbf{X}^* - E(\mathbf{X})\}\right]^2 \\ &= E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T E(\mathbf{X})\right\}^2 + E\{(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{M}(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})\}, \quad (33) \end{aligned}$$

For the first term in (33) notice that $\alpha = E(Y) - \sum_{j=1}^p \beta_j E(X_j)$ and $\tilde{\alpha} = \bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \bar{X}_j$; therefore,

$$E\left\{(\alpha - \tilde{\alpha}) + (\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T E(\mathbf{X})\right\}^2 = E\left[\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{\bar{X}_j - E(X_j)\} - E(Y)\right]^2,$$

and (32) is established. \square

Proof of Proposition 1.

Since for any numbers c, d, B we have that $\min(c + d, B) \leq \min(c, B) + \min(d, B)$ and $\min(c + d, B) \geq \min(c, B/2) + \min(d, B/2)$, then

$$\begin{aligned} R(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}) &= \lim_{B \rightarrow \infty} \lim_{n \rightarrow \infty} E \min\{n\tilde{L}(\tilde{\alpha}, \tilde{\boldsymbol{\beta}}), B\} = \lim_{B \rightarrow \infty} \lim_{n \rightarrow \infty} E \min\left(n\left[\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{\bar{X}_j - E(X_j)\} - E(Y)\right]^2, B\right) \\ &\quad + \lim_{B \rightarrow \infty} \lim_{n \rightarrow \infty} E \min\left(n\left\{(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})^T \mathbf{M}(\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}})\right\}, B\right). \quad (34) \end{aligned}$$

We start with the first summand in (34). Since $\sqrt{n}(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta})$ is asymptotically normal then

$$\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{\bar{X}_j - E(X_j)\} - E(Y) - \left[\bar{Y} - \sum_{j=1}^p \beta_j \{\bar{X}_j - E(X_j)\} - E(Y)\right] = \sum_{j=1}^p (\beta_j - \tilde{\beta}_j) \{\bar{X}_j - E(X_j)\} = o_p(1/\sqrt{n}).$$

Therefore, by the continuous mapping theorem and since the random variables are bounded, then for any B ,

$$\begin{aligned} \lim_{n \rightarrow \infty} E \min \left(n \left[\bar{Y} - \sum_{j=1}^p \tilde{\beta}_j \{ \bar{X}_j - E(X_j) \} - E(Y) \right]^2, B \right) \\ = \lim_{n \rightarrow \infty} E \min \left(n \left[\bar{Y} - \sum_{j=1}^p \beta_j \{ \bar{X}_j - E(X_j) \} - E(Y) \right]^2, B \right) \\ = E \min (Z_1^2, B), \end{aligned} \quad (35)$$

where Z_1 is normal with mean zero and variance $\left[Y - \sum_{j=1}^p \beta_j \{ X_j - E(X_j) \} - E(Y) \right]^2$.

For the second summand in (34), the asymptotic normality $\sqrt{n}(\tilde{\beta} - \beta) \xrightarrow{D} N(0, \Sigma)$ implies that for any B ,

$$\lim_{n \rightarrow \infty} E \min \left(n \left\{ (\beta - \tilde{\beta})^T \mathbf{M} (\beta - \tilde{\beta}) \right\}, B \right) = E \min \{ \mathbf{Z}_2^T \mathbf{Z}_2, B \}, \quad (36)$$

where \mathbf{Z}_2 is a normal vector with mean zero and variance matrix $\mathbf{M}^{1/2} \Sigma \mathbf{M}^{1/2}$. Taking limits as $B \rightarrow \infty$ in (35) and (36) completes the proof of the proposition since $\text{Trace}(\mathbf{M}^{1/2} \Sigma \mathbf{M}^{1/2}) = \text{Trace}(\mathbf{M} \Sigma)$. \square

We will not prove Theorem 1 since it is a special case of Theorem 2.

Proof of Theorem 2.

Part (i) We have that (Buja et al., 2017) $\{\hat{\beta}_{LSE}\}_j = \beta_j + \frac{\langle \delta, \mathbf{X}_{j\bullet} \rangle}{\|\mathbf{X}_{j\bullet}\|^2}$. Furthermore,

$$\{\hat{\beta}_{LSE}\}_j = \beta_j + \frac{\langle \delta, \mathbf{X}_{j\bullet} \rangle}{\|\mathbf{X}_{j\bullet}\|^2} = \beta_j + \frac{\frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)} \delta^{(i)}}{\frac{1}{n} \sum_{i=1}^n \{X_{j\bullet}^{(i)}\}^2} = \beta_j + \frac{\frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)} \delta^{(i)}}{E(X_{j\bullet}^2)} + o_p(1/\sqrt{n}),$$

since

$$\frac{\sqrt{n}}{n} \sum_{i=1}^n \delta^{(i)} (X_{j\bullet}^{(i)} - X_{j\bullet}) = \frac{\sqrt{n}}{n} \sum_{i=1}^n \delta^{(i)} \{ \bar{\mathbf{X}}_{-j}^{(i)} \}^T \left(\hat{\beta}_{-j\bullet} - \beta_{-j\bullet} \right) \xrightarrow{p} 0, \text{ and } \frac{1}{n} \|\mathbf{X}_{j\bullet}\|^2 \xrightarrow{p} E(X_{j\bullet}^2).$$

The results now follows from the multivariate CLT and Slutsky's theorem.

Part (ii) For each j ,

$$\begin{aligned} \{\hat{\beta}_{TI}\}_j &= \bar{W}_j - \hat{a}\bar{U}_1 - \sum_{j'=1}^p \hat{b}_{j'} \bar{U}_{j'+1} = \bar{W}_j - a\bar{U}_1 - \sum_{j'=1}^p b_{j'} \bar{U}_{j'+1} + o_p(1/\sqrt{n}) \\ &= \beta_j + \frac{1}{n} \sum_{i=1}^n \tilde{\delta}_j^{(i)} + o_p(1/\sqrt{n}). \end{aligned} \quad (37)$$

Again, the multivariate CLT and Slutsky's theorem imply the result.

To prove the furthermore part, notice that for each j ,

$$\begin{aligned} V_j &= \frac{\delta X_{j\bullet}}{E(X_{j\bullet}^2)} - \tilde{\delta}_j = \frac{\left(Y - \alpha - \sum_{j'=1}^p \beta_{j'} X_{j'} \right) X_{j\bullet}}{E(X_{j\bullet}^2)} - \frac{Y X_{j\bullet}}{E(X_{j\bullet}^2)} + \beta_j - b_j + a \frac{X_{j\bullet}}{E(X_{j\bullet}^2)} + \sum_{j'=1}^p b_{j'} \frac{X_{j'} X_{j\bullet}}{E(X_{j\bullet}^2)} \\ &= \beta_j - b_j + (a - \alpha) \frac{X_{j\bullet}}{E(X_{j\bullet}^2)} + \sum_{j'=1}^p (b_{j'} - \beta_{j'}) \frac{X_{j'} X_{j\bullet}}{E(X_{j\bullet}^2)} = (a - \alpha) U_1 + \sum_{j'=1}^p (b_{j'} - \beta_{j'}) U_{j'+1} \end{aligned} \quad (38)$$

is a linear function of U_1, \dots, U_{p+1} and hence is orthogonal to $\tilde{\delta}_j$, i.e., $EV_j\tilde{\delta}_j = 0$ for all j . Therefore,

$$\Sigma_{LSE} = Cov(\mathbf{X}_\bullet \delta) = Cov\left\{\tilde{\delta} + (\mathbf{X}_\bullet \delta - \tilde{\delta})\right\} = Cov(\tilde{\delta} + \mathbf{V}) = Cov(\tilde{\delta}) + Cov(\mathbf{V}) = \Sigma_{TI} + \Sigma_{\text{diff}}.$$

Part (iii) We show that

$$\{\hat{\beta}_{PI}\}_j - \{\hat{\beta}_{TI}\}_j = \frac{1}{n+m} \sum_{i=1}^{n+m} V_j^{(i)} + o_p(1/\sqrt{n}). \quad (39)$$

Hence, $\sqrt{n}(\hat{\beta}_{PI} - \hat{\beta}_{TI}) \xrightarrow{\mathcal{D}} N(0, \nu \Sigma_{\text{diff}})$. Since \mathbf{V} is orthogonal to $\tilde{\delta}$, then

$$\sqrt{n}(\hat{\beta}_{PI} - \beta) = \sqrt{n}(\hat{\beta}_{TI} - \beta) + \sqrt{n}(\hat{\beta}_{PI} - \hat{\beta}_{TI}) \xrightarrow{\mathcal{D}} N(0, \Sigma_{TI} + \nu \Sigma_{\text{diff}}).$$

We now prove (39) by a somewhat lengthy calculation. We have that

$$\{\hat{\beta}_{PI}\}_j = \bar{W}_j - \hat{a}\bar{U}_1 - \sum_{j'=1}^p \hat{b}_{j'} \bar{U}_{j'+1} = \bar{W}_j - a\bar{U}_1 - \sum_{j'=1}^p b_{j'} \bar{U}_{j'+1} + o_p(1/\sqrt{n}). \quad (40)$$

Therefore, (40) and (37) yield

$$\{\hat{\beta}_{PI}\}_j - \{\hat{\beta}_{TI}\}_j = \bar{W}_j - \bar{W}_j - a(\bar{U}_1 - \bar{U}_1) - \sum_{j'=1}^p b_{j'}(\bar{U}_{j'+1} - \bar{U}_{j'+1}) + o_p(1/\sqrt{n}).$$

We have that

$$\bar{W}_j - \bar{W}_j = \alpha(\bar{U}_1 - \bar{U}_1) + \sum_{j'=1}^p \beta_{j'}(\bar{U}_{j'+1} - \bar{U}_{j'+1}) + o_p(1/\sqrt{n}),$$

and hence,

$$\{\hat{\beta}_{PI}\}_j - \{\hat{\beta}_{TI}\}_j = (\alpha - a)(\bar{U}_1 - \bar{U}_1) + \sum_{j'=1}^p (\beta_{j'} - b_{j'}) (\bar{U}_{j'+1} - \bar{U}_{j'+1}) + o_p(1/\sqrt{n}). \quad (41)$$

We now consider the summands in (41). We start with the first summand

$$\begin{aligned} \bar{U}_1 - \bar{U} &= \frac{\frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)}}{\check{E}(X_{j\bullet}^2)} - \frac{\frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)}}{E(X_{j\bullet})^2} = \frac{1}{n} \sum_{i=1}^n X_{j\bullet}^{(i)} \left(\frac{1}{\check{E}(X_{j\bullet}^2)} - \frac{1}{E(X_{j\bullet})^2} \right) + \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n} \sum_{i=1}^n \bar{\mathbf{X}}_{-j}^{(i)}}{\check{E}(X_{j\bullet}^2)} \\ &= \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n} \sum_{i=1}^n \bar{\mathbf{X}}_{-j}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}) \\ &= \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} \bar{\mathbf{X}}_{-j}^{(i)}}{E(X_{j\bullet})^2} + \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \left(\frac{1}{n} \sum_{i=1}^n \bar{\mathbf{X}}_{-j}^{(i)} - \frac{1}{n+m} \sum_{i=1}^{n+m} \bar{\mathbf{X}}_{-j}^{(i)} \right)}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}) \\ &= \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} \bar{\mathbf{X}}_{-j}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}). \end{aligned}$$

Since

$$0 = \check{E}(X_{j\bullet}) = \check{E}(X_j - \beta_{-j\bullet}^T \bar{\mathbf{X}}_{-j}) \implies \check{E}(X_j) = \check{E}(\beta_{-j\bullet}^T \bar{\mathbf{X}}_{-j})$$

then,

$$\check{E}(X_{j\bullet}) = \check{E}(X_j - \beta_{-j\bullet}^T \bar{\mathbf{X}}_{-j}) = \check{E}(\beta_{-j\bullet}^T \bar{\mathbf{X}}_{-j}) - \check{E}(\beta_{-j\bullet}^T \bar{\mathbf{X}}_{-j}) = (\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} \bar{\mathbf{X}}_{-j}^{(i)},$$

and therefore,

$$\bar{U}_1 - \bar{U}_1 = \frac{-\frac{1}{n+m} \sum_{i=1}^{n+m} X_{j\bullet}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}). \quad (42)$$

We now consider the j' -th term in (41) for $j' \neq j$; then, $E(X_{j'} X_{j\bullet}) = 0$. We have

$$\begin{aligned} \bar{U}_{j'+1} - \bar{U}_{j'+1} &= \frac{\frac{1}{n} \sum_{i=1}^n X_{j'}^{(i)} X_{j\bullet}^{(i)}}{\check{E}(X_{j\bullet}^2)} - \frac{\frac{1}{n} \sum_{i=1}^n X_{j'}^{(i)} X_{j\bullet}^{(i)}}{E(X_{j\bullet})^2} \\ &= \frac{1}{n} \sum_{i=1}^n X_{j'}^{(i)} X_{j\bullet}^{(i)} \left\{ \frac{1}{\check{E}(X_{j\bullet}^2)} - \frac{1}{E(X_{j\bullet})^2} \right\} + \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n} \sum_{i=1}^n X_{j'}^{(i)} \vec{\mathbf{X}}_{-j}^{(i)}}{\check{E}(X_{j\bullet}^2)} \\ &= \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} X_{j'}^{(i)} \vec{\mathbf{X}}_{-j}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}). \end{aligned}$$

Since

$$0 = \check{E}(X_{j\bullet} X_{j'}) = \check{E}\{(X_j - \beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j}) X_{j'}\} \implies \check{E}(X_j X_{j'}) = \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_{j'})$$

then,

$$\check{E}(X_{j\bullet} X_{j'}) = \check{E}\{(X_j - \beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j}) X_{j'}\} = \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_{j'}) - \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_{j'}) = (\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} X_{j'}^{(i)} \vec{\mathbf{X}}_{-j}^{(i)},$$

and therefore,

$$\bar{U}_{j'+1} - \bar{U}_{j'+1} = \frac{-\frac{1}{n+m} \sum_{i=1}^{n+m} X_{j\bullet}^{(i)} X_{j'}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}). \quad (43)$$

For the j -th term in (41), we have that $E(X_j X_{j\bullet}) = E(X_{j\bullet})^2$; hence,

$$\begin{aligned} \bar{U}_{j+1} - \bar{U}_{j+1} &= \frac{\frac{1}{n} \sum_{i=1}^n X_j^{(i)} X_{j\bullet}^{(i)}}{\check{E}(X_{j\bullet}^2)} - \frac{\frac{1}{n} \sum_{i=1}^n X_j^{(i)} X_{j\bullet}^{(i)}}{E(X_{j\bullet})^2} \\ &= \left\{ \frac{1}{n} \sum_{i=1}^n X_j^{(i)} X_{j\bullet}^{(i)} - E(X_{j\bullet})^2 \right\} \left\{ \frac{1}{\check{E}(X_{j\bullet}^2)} - \frac{1}{E(X_{j\bullet})^2} \right\} + \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n} \sum_{i=1}^n X_j^{(i)} \vec{\mathbf{X}}_{-j}^{(i)} + E(X_{j\bullet})^2 - \check{E}(X_{j\bullet}^2)}{\check{E}(X_{j\bullet}^2)} \\ &= \frac{(\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n} \sum_{i=1}^n X_j^{(i)} \vec{\mathbf{X}}_{-j}^{(i)} + E(X_{j\bullet})^2 - \check{E}(X_{j\bullet}^2)}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}). \end{aligned}$$

Since

$$\check{E}(X_{j\bullet}^2) = \check{E}(X_{j\bullet} X_j) = \check{E}\{(X_j - \beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j}) X_j\} \implies \check{E}(X_j^2) = \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_j) + \check{E}(X_{j\bullet}^2)$$

then,

$$\begin{aligned} \check{E}(X_{j\bullet} X_j) &= \check{E}\{(X_j - \beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j}) X_j\} = \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_j) + \check{E}(X_{j\bullet}^2) - \check{E}(\beta_{-j\bullet}^T \vec{\mathbf{X}}_{-j} X_j) \\ &= (\beta_{-j\bullet} - \beta_{-j\bullet})^T \frac{1}{n+m} \sum_{i=1}^{n+m} X_j^{(i)} \vec{\mathbf{X}}_{-j}^{(i)} + \check{E}(X_{j\bullet}^2). \end{aligned}$$

Thus, for the j -th term we obtain,

$$\bar{U}_{j+1} - \bar{U}_{j+1} = \frac{-\frac{1}{n+m} \sum_{i=1}^{n+m} X_{j\bullet}^{(i)} X_j^{(i)}}{E(X_{j\bullet})^2} + 1 + o_p(1/\sqrt{n}). \quad (44)$$

Tracking back to (41) through (42),(43), and (44), we obtain

$$\{\hat{\beta}_{PI}\}_j - \{\hat{\beta}_{TI}\}_j = (\beta_j - b_j) + (a - \alpha) \frac{\frac{1}{n+m} \sum_{i=1}^{n+m} X_{j\bullet}^{(i)}}{E(X_{j\bullet})^2} + \sum_{j'=1}^p (b_{j'} - \beta_{j'}) \frac{\frac{1}{n+m} \sum_{i=1}^{n+m} X_{j\bullet}^{(i)} X_{j'}^{(i)}}{E(X_{j\bullet})^2} + o_p(1/\sqrt{n}).$$

Hence, (38) implies that

$$\{\hat{\beta}_{PI}\}_j - \{\hat{\beta}_{TI}\}_j = \frac{1}{n+m} \sum_{i=1}^{n+m} V_j^{(i)} + o_p(1/\sqrt{n}),$$

and (39) is established. \square

Proof of Theorem 3. (Sufficiency). The same proof outlined in the introduction applies. We bring it here for completeness. We have that

$$\delta X - \tilde{\delta}_K = \sum_{k=1}^K b_k q_k(X) - \alpha X - \beta(X^2 - 1).$$

If $X \in \text{span}\{q_1(X), \dots, q_K(X)\}$ and $X^2 - 1 \in \text{span}\{q_1(X), \dots, q_K(X)\}$, then $\delta X - \tilde{\delta}_K$ is orthogonal to $\tilde{\delta}_K$ and then

$$E_G(\delta X)^2 = E_G\left(\tilde{\delta}_K + \delta X - \tilde{\delta}_K\right)^2 = E_G \tilde{\delta}_K^2 + E_G\left(\delta X - \tilde{\delta}_K\right)^2 \geq E_G \tilde{\delta}_K^2.$$

(Necessity). Let G be such that X and Y are linearly dependent, i.e., $Y = \alpha + \beta X$ for certain α and β . Then $\delta \equiv 0$. We have that

$$XY = \beta + \alpha X + \beta(X^2 - 1)$$

and therefore $\tilde{\delta}_K \equiv 0$ in Model (24) for every α, β iff $X \in \text{span}\{q_1(X), \dots, q_K(X)\}$ and $X^2 - 1 \in \text{span}\{q_1(X), \dots, q_K(X)\}$. \square

References

- Brown, L.D. (1990). An Ancillarity Paradox Which Appears in Multiple Linear Regression. *Annals of Statistics*, **18**, 471–493.
- Bühlmann, P., van de Geer, S. (2015). High-dimensional inference in misspecified linear models. *Electronic Journal of Statistics*, **9**, 1449–1473.
- Buja, A., Berk, R., Brown, L., George, E., Pitkin, E., Traskin, M., Zhao, L., Zhang, K. (2017). Models as Approximations, Part I: A Conspiracy of Nonlinearity and Random Regressors in Linear Regression. *Submitted*; arXiv:1404.1578.

- Goldberg, A., Zhu, X., Singh, A., Xu, Z., Nowak, R. (2009). Multi-manifold semi-supervised learning, in Proceedings of the 12th International Conference on Artificial Intelligence and Statistics, 2009, pp. 1691-176.
- Chakraborty A., Cai T. (2018). Efficient and Adaptive Linear Regression in Semi-Supervised settings. *Annals of Statistics*, **46**, 1541–1572.
- Cochran W.G. (1977). *Sampling Techniques* (3rd ed.), New York: John Wiley.
- Johnson, R., Zhang, T. (2008). Graph-based semi-supervised learning and spectral kernel design. *IEEE Transactions on Information Theory*. **54**, 275–288.
- Kriegler, B., Berk, R. (2010). Small area estimation of the homeless in Los Angeles: An application of cost-sensitive stochastic gradient boosting. *Annals of Applied Statistics*, **4**, 1234–1255.
- Lafferty, J., Wasserman, L. (2008). Statistical analysis of semi-supervised regression. *Advances in Neural Information Processing Systems*, **20**, 801-808.
- Liang, F., Mukherjee, S., West, M. (2007). The use of unlabeled data in predictive modeling. *Statistical Science*, **22**, 189-205.
- Le Cam, L., Yang, G.L. (1990). *Asymptotics in Statistics : Some Basic Concepts*. New York: Springer.
- Little, R.J.A. (1992). Regression With Missing X's: A Review. *Journal of the American Statistical Association*, **87**, 1227–1237.
- Merz, C.J., Clair, D.C.St., Bond, W.E (1992). Semi-supervised adaptive resonance theory (smart2). *International Joint Conference on Neural Networks*, **3**, 851-856.
- Wang, J., Shen, X., Pan, W. (2009). On efficient large margin semisupervised learning: Method and theory. *The Journal of Machine Learning Research*, **10**, 719-742.
- White, H. (1980). Using Least Squares to Approximate Unknown Regression Functions. *International Economic Review*, **21**, 149–170.
- Zhang, A., Brown L.D., Cai, T. (2017). Semi-supervised Inference: General Theory and estimation of means. Submitted; arXiv:1606.07268.
- Zhou, Z.H., Li, M. (2005). Semi-supervised regression with co-training. *International Joint Conference on Artificial Intelligence*, **5**, 908–913.

Zhu, X., (2005). Semi-Supervised Learning Literature Survey. Computer Sciences, University of Wisconsin-Madison, **1530**.