

EFFICIENT AND ADAPTIVE LINEAR REGRESSION IN SEMI-SUPERVISED SETTINGS

BY ABHISHEK CHAKRABORTTY AND TIANXI CAI*

Harvard University

We consider the linear regression problem under semi-supervised settings wherein the available data typically consists of: (i) a small or moderate sized ‘labeled’ data, and (ii) a *much larger* sized ‘unlabeled’ data. Such data arises naturally from settings where the outcome, unlike the covariates, is expensive to obtain, a frequent scenario in modern studies involving large databases like electronic medical records (EMR). Supervised estimators like the ordinary least squares (OLS) estimator utilize only the labeled data. It is often of interest to investigate if and when the unlabeled data can be exploited to improve estimation of the regression parameter in the adopted linear model.

In this paper, we propose a class of ‘Efficient and Adaptive Semi-Supervised Estimators’ (EASE) to improve estimation efficiency. The EASE are two-step estimators adaptive to model mis-specification, leading to improved (optimal in some cases) efficiency under model mis-specification, and equal (optimal) efficiency under a linear model. This adaptive property, often unaddressed in the existing literature, is crucial for advocating ‘safe’ use of the unlabeled data. The construction of EASE primarily involves a flexible ‘semi-non-parametric’ imputation, including a smoothing step that works well even when the number of covariates is not small; and a follow up ‘refitting’ step along with a cross-validation (CV) strategy both of which have useful practical as well as theoretical implications towards addressing two important issues: under-smoothing and over-fitting. We establish asymptotic results including consistency, asymptotic normality and the adaptive properties of EASE. We also provide influence function expansions and a ‘double’ CV strategy for inference. The results are further validated through extensive simulations, followed by application to an EMR study on auto-immunity.

1. Introduction. In recent years, semi-supervised learning (SSL) has emerged as an exciting new area of research in statistics and machine learning. A detailed discussion on SSL including its practical relevance, the primary question of interest in SSL, and the existing relevant literature can be found in [Chapelle, Schölkopf and Zien \(2006\)](#) and [Zhu \(2008\)](#). A typi-

*This research was partially supported by National Institutes of Health grants R01 HL089778, U54 HG007963, and U01 HL121518.

Keywords and phrases: Semi-supervised Linear Regression, Semi-parametric Inference, Model Mis-specification, Adaptive Estimation, Semi-non-parametric Imputation.

cal semi-supervised (SS) setting is characterized by two types of available data: (i) a small or moderate sized ‘labeled’ data, \mathcal{L} , containing observations for both an outcome Y and a set of covariates \mathbf{X} of interest, and (ii) an ‘unlabeled’ data, \mathcal{U} , of *much larger* size but having observations *only* for the covariates \mathbf{X} . By virtue of its large size, \mathcal{U} essentially gives us the distribution of \mathbf{X} , denoted henceforth by $\mathbb{P}_{\mathbf{X}}$. Such a setting arises naturally whenever the covariates are easily available so that unlabeled data is plentiful, but the outcome is costly or difficult to obtain, thereby limiting the size of \mathcal{L} . This scenario is directly relevant to a variety of practical problems, especially in the modern ‘big data’ era, with massive unlabeled datasets (often electronically recorded) becoming increasingly available and tractable. A few familiar examples include machine learning problems like text mining, web page classification, speech recognition, natural language processing etc.

Among biomedical applications, a particularly interesting problem where SSL can be of great use is the statistical analysis of electronic medical records (EMR) data. Endowed with a wealth of de-identified clinical and phenotype data for large patient cohorts, EMR linked with bio-repositories are increasingly gaining popularity as rich resources of data for discovery research (Kohane, 2011). Such large scale datasets obtained in a cost-effective and timely manner are of great importance in modern medical research for addressing important questions such as the biological role of genetic variants in disease susceptibility and progression (Kohane, 2011). However, one major bottleneck impeding EMR driven research is the difficulty in obtaining validated phenotype information (Liao et al., 2010) since they are labor intensive or expensive to obtain. Thus, gold standard labels and genomic measurements are typically available only for a small subset nested within a large cohort. In contrast, digitally recorded data on the clinical variables are often available on all subjects, highlighting the necessity and utility of developing robust SSL methods that can leverage such rich source of auxiliary information to improve phenotype definition and estimation precision.

SSL primarily distinguishes from standard supervised methods by making use of \mathcal{U} , an information that is ignored by the latter. The ultimate question of interest in SSL is to investigate if and when the information on $\mathbb{P}_{\mathbf{X}}$ in \mathcal{U} can be exploited to improve the efficiency over a given supervised approach. In recent years, several graph based non-parametric SSL approaches have been proposed (Zhu, 2005; Belkin, Niyogi and Sindhwani, 2006) for regression or classification. These approaches essentially target non-parametric SS estimation of $\mathbb{E}(Y|\mathbf{X})$ and therefore, for provable improvement guarantees, must rely implicitly or explicitly on assumptions relating $\mathbb{P}_{\mathbf{X}}$ to $\mathbb{P}_{Y|\mathbf{X}}$ (the conditional distribution of Y given \mathbf{X}), as duly noted and characterized more

formally in [Lafferty and Wasserman \(2007\)](#). For non-parametric classification problems, the theoretical underpinnings of SSL including its scope and the consequences of using \mathcal{U} have been also studied earlier by [Castelli and Cover \(1995, 1996\)](#). More parametric SS approaches, still aimed mostly at prediction, have also been studied for classification, including the ‘generative model’ approach ([Nigam et al., 2000](#); [Nigam, 2001](#)) which is based on modeling the joint distribution of (Y, \mathbf{X}) as an identifiable mixture of parametric models, thereby implicitly relating $\mathbb{P}_{Y|\mathbf{X}}$ and $\mathbb{P}_{\mathbf{X}}$. However, these approaches depend strongly on the validity of the assumed mixture model, violation of which can actually *degrade* their performance compared to the supervised approach ([Cozman and Cohen, 2001](#); [Cozman, Cohen and Cirelo, 2003](#)).

However SS *estimation problems*, especially from a semi-parametric point of view, has been somewhat less studied in SSL. Such problems are generally aimed at estimating some (finite-dimensional) parameter $\theta_0 \equiv \theta_0(\mathbb{P})$, where $\mathbb{P} = (\mathbb{P}_{Y|\mathbf{X}}, \mathbb{P}_{\mathbf{X}})$, and the key to the potential usefulness of \mathcal{U} in improving estimation of θ_0 lies in understanding when $\theta_0(\mathbb{P})$ relates to $\mathbb{P}_{\mathbf{X}}$. For simple parameters like $\theta_0(\mathbb{P}) = \mathbb{E}(Y)$, unless $\mathbb{E}(Y|\mathbf{X})$ is a constant, θ_0 clearly depends on $\mathbb{P}_{\mathbf{X}}$ and hence, improved SS estimation is possible compared to the supervised estimator $\bar{Y}_{\mathcal{L}}$, the sample mean of Y based on \mathcal{L} . The situation is however more subtle for other choices of θ_0 , especially those where θ_0 is the target parameter corresponding to an underlying parametric *working* model for $\mathbb{P}_{Y|\mathbf{X}}$. This includes the least squares parameter, as studied in this paper, targeted by a working linear model for $\mathbb{E}(Y|\mathbf{X})$. Such models are often adopted due to their appealing simplicity and interpretability.

In general, for such cases, if the adopted working model for $\mathbb{P}_{Y|\mathbf{X}}$ is correct and θ_0 is not related to $\mathbb{P}_{\mathbf{X}}$, then one *cannot* possibly gain through SSL by using the knowledge of $\mathbb{P}_{\mathbf{X}}$ ([Zhang and Oles, 2000](#); [Seeger, 2002](#)). On the other hand, under model mis-specification, θ_0 may inherently *depend* on $\mathbb{P}_{\mathbf{X}}$, and thus imply the potential utility of \mathcal{U} in improving the estimation. However, inappropriate use of \mathcal{U} may lead to degradation of the estimation precision. This therefore signifies the need for *robust* and efficient SS estimators that are *adaptive* to model mis-specification, so that they are as efficient as the supervised estimator under the correct model and more efficient under model mis-specification. To the best of our knowledge, work done along these lines is relatively scarce in the SSL literature, one notable exception being the recent work of [Kawakita and Kanamori \(2013\)](#), where they use a very different approach based on density ratio estimation, building on the more restrictive approach of [Sokolovska, Cappé and Yvon \(2008\)](#). However, as we observe in our simulation studies, the extent of the efficiency gain actually achieved by these approaches can be quite incremental, at least in finite

samples. Further, the seemingly unclear choice of the ideal (nuisance) model to be used for density ratio estimation can also have a significant impact on the performance, both finite sample and asymptotic, of these estimators.

We propose here a class of Efficient and Adaptive Semi-Supervised Estimators (EASE) in the context of linear regression problems. We essentially adopt a semi-parametric perspective wherein the adopted linear ‘working’ model can be potentially mis-specified, and the goal is to obtain efficient and adaptive SS estimators of the regression parameter through robust usage of \mathcal{U} . The EASE are two-step estimators with a simple and scalable construction based on a first step of ‘semi-non-parametric’ (SNP) imputation which includes a smoothing step and a follow-up ‘refitting’ step. In the second step, we regress the imputed outcomes against the covariates using the unlabeled data to obtain our SNP imputation based SS estimator, and then further combine it optimally with the supervised estimator to obtain the final EASE. Dimension reduction methods are also employed in the smoothing step to accommodate higher dimensional \mathbf{X} , if necessary. Further, we extensively adopt cross-validation (CV) techniques in the imputation, leading to some useful theoretical properties (apart from practical benefits) typically not observed for smoothing based two-step estimators. We demonstrate that EASE is guaranteed to be efficient and adaptive in the sense discussed above, and also achieves semi-parametric optimality whenever the SNP imputation is ‘sufficient’ or the linear model holds. We also provide data adaptive methods to optimally select the directions for smoothing when dimension reduction is desired, and tools for inference with EASE.

The rest of this paper is organized as follows. In Section 2, we formulate the SS linear regression problem. In Section 3, we construct a family of SS estimators based on SNP imputation and establish all their properties, and further propose the EASE as a refinement of these estimators. For all our proposed estimators, we also address their associated inference procedures based on ‘double’ CV methods. In Section 4, we discuss a kernel smoothing based implementation of the SNP imputation and establish all its properties. In Section 5, we discuss SS dimension reduction techniques, useful for implementing the SNP imputation. Simulation results and an application to an EMR study are shown in Section 6, followed by concluding discussions in Section 7. Proofs of all theoretical results, and further numerical results and discussions are distributed in the [Appendix](#) and [Supplementary Material](#).

2. Problem Set-up.

Data Representation. Let $Y \in \mathbb{R}$ denote the outcome random variable and $\mathbf{X} \in \mathbb{R}^p$ denote the covariate vector, where p is fixed, and let $\mathbf{Z} = (Y, \mathbf{X}')$.

Then the entire data available for analysis can be represented as $\mathbb{S} = (\mathcal{L} \cup \mathcal{U})$, where $\mathcal{L} = \{\mathbf{Z}_i \equiv (Y_i, \mathbf{X}_i')' : i = 1, \dots, n\}$ consists of n independent and identically distributed (i.i.d.) observations from the joint distribution $\mathbb{P}_{\mathbf{Z}}$ of \mathbf{Z} , $\mathcal{U} = \{\mathbf{X}_i : i = n+1, \dots, n+N\}$ consists of N i.i.d. observations from $\mathbb{P}_{\mathbf{X}}$, and $\mathcal{L} \perp \mathcal{U}$. Throughout, for notational convenience, we use the subscript ‘ j ’ to denote the unlabeled observations, and re-index without loss of generality (w.l.o.g.) the N observations in \mathcal{U} as: $\mathcal{U} = \{\mathbf{X}_j : j = n+1, \dots, n+N\}$.

ASSUMPTION 2.1 (Basic Assumptions). (a) We assume that \mathbf{Z} has finite 2^{nd} moments and $\boldsymbol{\Sigma} \equiv \text{Var}(\mathbf{X})$ is positive definite, denoted as $\boldsymbol{\Sigma} \succ 0$. We also assume, for simplicity, that \mathbf{X} has a compact support $\mathcal{X} \subseteq \mathbb{R}^p$.

(b) We assume $N \gg n$ i.e. $n/N \rightarrow 0$ as $n, N \rightarrow \infty$, and \mathcal{L} and \mathcal{U} arise from the same underlying distribution, i.e. $\mathbf{Z} \sim \mathbb{P}_{\mathbf{Z}}$ for all subjects in \mathcal{S} .

Notations. Let $\boldsymbol{\Gamma} = \mathbb{E}(\overrightarrow{\mathbf{X}} \overrightarrow{\mathbf{X}}') \succ 0$, where $\forall \mathbf{v} \in \mathbb{R}^p$, $\overrightarrow{\mathbf{v}} = (1, \mathbf{v}')' \in \mathbb{R}^{(p+1)}$. Let $\mathcal{L}_2(\mathbb{P}_{\mathbf{X}})$ denote the space of all \mathbb{R} -valued measurable functions of \mathbf{X} having finite L_2 norm with respect to (w.r.t.) $\mathbb{P}_{\mathbf{X}}$, and for any $g(\cdot) \in \mathcal{L}_2(\mathbb{P}_{\mathbf{X}})$, let $\boldsymbol{\Sigma}(g) \succ 0$ denote the $(p+1) \times (p+1)$ matrix $\boldsymbol{\Gamma}^{-1} \mathbb{E}[\overrightarrow{\mathbf{X}} \overrightarrow{\mathbf{X}}' \{Y - g(\mathbf{X})\}^2] \boldsymbol{\Gamma}^{-1}$. Lastly, let $\|\cdot\|$ denote the L_2 vector norm, and for any integer $a \geq 1$, let I_a denote the identity matrix of order a , and $\mathcal{N}_a[\boldsymbol{\mu}, \boldsymbol{\Omega}]$ denote the a -variate Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^a$ and covariance matrix $\boldsymbol{\Omega}_{a \times a} \succ 0$.

REMARK 2.1. Assumption 2.1 (b) enlists some fundamental characteristics of SS settings. Indeed, the condition of \mathcal{L} and \mathcal{U} being equally distributed has usually been an integral part of the *definition* of SS settings (Chapelle, Schölkopf and Zien, 2006; Kawakita and Kanamori, 2013). Interpreted in missing data terminology, it entails that Y in \mathcal{U} are ‘missing completely at random’ (MCAR), with the missingness/labeling being typically by design. Interestingly, the crucial assumption of MCAR, although commonly required, has often stayed implicit in the SSL literature (Lafferty and Wasserman, 2007). It is important to note that while the SS set-up can be viewed as a missing data problem, it is quite *different* from standard ones, since with $n/N \rightarrow 0$ i.e. $|\mathcal{U}| \gg |\mathcal{L}|$, the proportion of Y observed in \mathcal{S} tends to 0 in SSL. Hence, the ‘positivity assumption’ typical in missing data theory, requiring this proportion to be bounded away from 0, is violated here. It is also worth noting that owing to such violations, the analysis of SS settings under more general missingness mechanisms such as ‘missing at random’ (MAR) is considerably more complicated and to our knowledge, the literature for SS *estimation* problems under such settings is virtually non-existent. Furthermore, for such problems, the traditional goal in SSL, that of improving upon a ‘supervised’ estimator, can become unclear without

MCAR, unless an appropriately weighted version of the supervised estimator is considered. Given these subtleties and the traditional assumptions (often implicit) in SSL, the MCAR condition is assumed for most of this paper, although a brief discussion on possible extensions of our proposed SS estimators to MAR settings is provided in the [Supplementary Material](#).

2.1. *The Target Parameter and Its Supervised Estimator.* We consider the linear regression *working model* given by:

$$(2.1) \quad Y = \vec{\mathbf{X}}'\boldsymbol{\theta} + \epsilon, \quad \text{with} \quad \mathbb{E}(\epsilon | \mathbf{X}) = 0,$$

where, $\boldsymbol{\theta} \in \mathbb{R}^{(p+1)}$ is an unknown regression parameter. Accounting for the potential mis-specification of the working model (2.1), we define the target parameter of interest as a model free parameter, as follows:

DEFINITION 2.1. The target parameter $\boldsymbol{\theta}_0$ for linear regression may be defined as the solution to the normal equations: $\mathbb{E}\{\vec{\mathbf{X}}(Y - \vec{\mathbf{X}}'\boldsymbol{\theta})\} = \mathbf{0}$ in $\boldsymbol{\theta} \in \mathbb{R}^{(p+1)}$, or equivalently, $\boldsymbol{\theta}_0 = \underset{\boldsymbol{\theta} \in \mathbb{R}^{(p+1)}}{\operatorname{argmin}} \mathbb{E}(Y - \vec{\mathbf{X}}'\boldsymbol{\theta})^2$.

Existence and uniqueness of $\boldsymbol{\theta}_0$ in 2.1 is clear. Further, $\vec{\mathbf{X}}'\boldsymbol{\theta}_0$ is the L_2 projection of $\mathbb{E}(Y|\mathbf{X}) \in \mathcal{L}_2(\mathbb{P}_{\mathbf{X}})$ onto the subspace of all linear functions of \mathbf{X} and hence, is the best linear predictor of Y given \mathbf{X} . The linear model (2.1) is *correct* (else, *mis-specified*) if and only if $\mathbb{E}(Y|\mathbf{X})$ lies in this space (in which case, $\mathbb{E}(Y|\mathbf{X}) = \vec{\mathbf{X}}'\boldsymbol{\theta}_0$). When the model is correct, $\boldsymbol{\theta}_0$ depends only on $\mathbb{P}_{Y|\mathbf{X}}$, not on $\mathbb{P}_{\mathbf{X}}$. Hence, improved estimation of $\boldsymbol{\theta}_0$ through SSL is impossible in this case unless further assumptions relating $\boldsymbol{\theta}_0$ to $\mathbb{P}_{\mathbf{X}}$ are made. On the other hand, under model mis-specification, the normal equations defining $\boldsymbol{\theta}_0$ inherently depend on $\mathbb{P}_{\mathbf{X}}$, thereby implying the potential utility of SSL in improving the estimation of $\boldsymbol{\theta}_0$ in this case.

The usual supervised estimator of $\boldsymbol{\theta}_0$ is the OLS estimator $\hat{\boldsymbol{\theta}}$, the solution in $\boldsymbol{\theta}$ to the equation: $n^{-1} \sum_{i=1}^n \vec{\mathbf{X}}_i(Y_i - \vec{\mathbf{X}}_i'\boldsymbol{\theta}) = \mathbf{0}$, the normal equations based on \mathcal{L} . Under Assumption 2.1 (a), it is well known that as $n \rightarrow \infty$,

$$(2.2) \quad n^{\frac{1}{2}}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}_0(\mathbf{Z}_i) + O_p\left(n^{-\frac{1}{2}}\right) \xrightarrow{d} \mathcal{N}_{(p+1)}[\mathbf{0}, \boldsymbol{\Sigma}(g_{\boldsymbol{\theta}_0})],$$

where $\boldsymbol{\psi}_0(\mathbf{Z}) = \boldsymbol{\Gamma}^{-1}\{\vec{\mathbf{X}}(Y - \vec{\mathbf{X}}'\boldsymbol{\theta}_0)\}$ and $g_{\boldsymbol{\theta}}(\mathbf{X}) = \vec{\mathbf{X}}'\boldsymbol{\theta} \ \forall \boldsymbol{\theta} \in \mathbb{R}^{(p+1)}$.

Our primary goal is to obtain efficient SS estimators of $\boldsymbol{\theta}_0$ using the *entire* data \mathcal{S} and compare their efficiencies to that of $\hat{\boldsymbol{\theta}}$. It is worth noting that the estimation efficiency of $\boldsymbol{\theta}_0$ also relates to the predictive performance of the fitted linear model since its out-of-sample prediction error is directly related to the mean squared error (w.r.t. the $\boldsymbol{\Sigma}$ metric) of the parameter estimate.

3. A Family of Imputation Based Semi-Supervised Estimators.

If Y in \mathcal{U} were actually observed, then one would simply fit the working model to the entire data in \mathcal{S} for estimating θ_0 . Our general approach is precisely motivated by this intuition. We first attempt to impute the missing Y in \mathcal{U} based on suitable training of \mathcal{L} in step (I). Then in step (II), we fit the linear model (2.1) to \mathcal{U} with the imputed outcomes. Clearly, the imputation is critical. Inaccurate imputation would lead to biased estimate of θ_0 , while inadequate imputation would result in loss of efficiency. We next consider SS estimators constructed under two imputation strategies for step (I) including a fully non-parametric imputation based on kernel smoothing (KS), and a semi-non-parametric (SNP) imputation that involves a smoothing step and a follow up ‘refitting’ step. Although the construction of the final EASE is based on the SNP imputation strategy, it is helpful to begin with a discussion of the first strategy in order to appropriately motivate and elucidate the discussion on EASE and the SNP imputation strategy.

3.1. *A Simple SS Estimator via Fully Non-Parametric Imputation.* We present here an estimator based on a fully non-parametric imputation involving KS when p is small. For simplicity, we shall assume here that \mathbf{X} is continuous with a density $f(\cdot)$. Let $m(\mathbf{x}) = \mathbb{E}(Y | \mathbf{X} = \mathbf{x})$ and $l(\mathbf{x}) = m(\mathbf{x})f(\mathbf{x})$. Consider the local constant KS estimator of $m(\mathbf{x})$,

$$(3.1) \quad \hat{m}(\mathbf{x}) = \frac{\frac{1}{nh^p} \sum_{i=1}^n \{K_h(\mathbf{X}_i, \mathbf{x})\} Y_i}{\frac{1}{nh^p} \sum_{i=1}^n K_h(\mathbf{X}_i, \mathbf{x})} = \frac{\hat{l}(\mathbf{x})}{\hat{f}(\mathbf{x})},$$

where $K_h(\mathbf{u}, \mathbf{v}) = K\{(\mathbf{u}-\mathbf{v})/h\}$ with $K : \mathbb{R}^p \rightarrow \mathbb{R}$ being some suitable kernel function and $h = h(n) > 0$ being the bandwidth. With $\hat{m}(\cdot)$ as defined in (3.1), we now fit (2.1) to the imputed unlabeled data: $\{[\hat{m}(\mathbf{X}_j), \mathbf{X}'_j] : j = n+1, \dots, n+N\}$ and obtain a SS estimator $\hat{\theta}_{np}$ of θ_0 as the solution in θ to:

$$(3.2) \quad \frac{1}{N} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \{ \hat{m}(\mathbf{X}_j) - \vec{\mathbf{X}}_j' \theta \} = \mathbf{0}.$$

Here and throughout in our constructions of SS estimators, \mathcal{L} with either the true or the imputed Y is not included in the final fitting step mostly due to technical convenience in the asymptotic analysis of our estimators, and also due to the fact that the contribution of \mathcal{L} , included in any form, in the final fitting step is asymptotically negligible since $n/N \rightarrow 0$. In order to study the properties of $\hat{\theta}_{np}$, we would require uniform (in L_∞ norm) convergence of $\hat{m}(\cdot)$ to $m(\cdot)$, a problem that has been extensively studied in the non-parametric statistics literature (Newey, 1994; Andrews, 1995; Masry, 1996;

(Hansen, 2008) under fairly general settings and assumptions. In particular, we would assume the following regularity conditions to hold:

ASSUMPTION 3.1. (i) $K(\cdot)$ is a symmetric q^{th} order kernel for some integer $q \geq 2$. (ii) $K(\cdot)$ is bounded, Lipschitz continuous and has a bounded support $\mathcal{K} \subseteq \mathbb{R}^p$. (iii) $\mathbb{E}(|Y|^s) < \infty$ for some $s > 2$. $\mathbb{E}(|Y|^s | \mathbf{X} = \mathbf{x})f(\mathbf{x})$ and $f(\mathbf{x})$ are bounded on \mathcal{X} . (iv) $f(\mathbf{x})$ is bounded away from 0 on \mathcal{X} . (v) $m(\cdot)$ and $f(\cdot)$ are q times continuously differentiable with bounded q^{th} derivatives on some open set $\mathcal{X}_0 \supseteq \mathcal{X}$. (vi) For any $\delta > 0$, let $A_\delta \subseteq \mathbb{R}^p$ denote the set $\{(\mathbf{x} - \mathbf{X})/\delta : \mathbf{x} \in \mathcal{X}\}$. Then, for small enough δ , $A_\delta \supseteq \mathcal{K}$ almost surely (a.s.).

Conditions (i)-(v) are fairly standard in the literature. In (v), the set \mathcal{X}_0 is needed mostly to make the notion of differentiability well-defined, with both $m(\cdot)$ and $f(\cdot)$ understood to have been analytically extended over $(\mathcal{X}_0 \setminus \mathcal{X})$. Condition (vi) implicitly controls the tail behaviour of \mathbf{X} , requiring that perturbations of \mathbf{X} in the form of $(\mathbf{X} + \delta\phi)$ with $\phi \in \mathcal{K}$ (bounded) and δ small enough, belong to \mathcal{X} a.s. $[\mathbb{P}_{\mathbf{X}}]$. We now present our result on $\hat{\boldsymbol{\theta}}_{np}$.

THEOREM 3.1. Suppose $n^{\frac{1}{2}}h^q \rightarrow 0$ and $(\log n)/(n^{\frac{1}{2}}h^p) \rightarrow 0$ as $n \rightarrow \infty$, and let $r_n = n^{\frac{1}{2}}h^q + (\log n)/(n^{\frac{1}{2}}h^p) + (n/N)^{\frac{1}{2}}$. Then, under Assumption 3.1,

$$(3.3) \quad n^{\frac{1}{2}} \left(\hat{\boldsymbol{\theta}}_{np} - \boldsymbol{\theta}_0 \right) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}_{\text{eff}}(\mathbf{Z}_i) + O_p(r_n) \xrightarrow{d} \mathcal{N}_{(p+1)}[\mathbf{0}, \boldsymbol{\Sigma}(m)],$$

where $\boldsymbol{\psi}_{\text{eff}}(\mathbf{Z}) = \boldsymbol{\Gamma}^{-1}[\vec{\mathbf{X}}\{Y - m(\mathbf{X})\}]$.

REMARK 3.1. Theorem 3.1 establishes the efficient and adaptive nature of $\hat{\boldsymbol{\theta}}_{np}$. The asymptotic variance $\boldsymbol{\Sigma}(m)$ of $\hat{\boldsymbol{\theta}}_{np}$ satisfies $\boldsymbol{\Sigma}(g) - \boldsymbol{\Sigma}(m) \succeq 0 \forall g(\cdot) \in \mathcal{L}^2(\mathbf{X})$ and the inequality is strict unless $g(\cdot) = m(\cdot)$ a.s. $[\mathbb{P}_{\mathbf{X}}]$. Hence, $\hat{\boldsymbol{\theta}}_{np}$ is asymptotically *optimal* among the class of all regular and asymptotically linear (RAL) estimators of $\boldsymbol{\theta}_0$ with influence function (IF) of the form: $\boldsymbol{\Gamma}^{-1}[\vec{\mathbf{X}}\{Y - g(\mathbf{X})\}]$ with $g(\cdot) \in \mathcal{L}_2(\mathbb{P}_{\mathbf{X}})$. In particular, $\hat{\boldsymbol{\theta}}_{np}$ is more efficient than $\hat{\boldsymbol{\theta}}$ whenever (2.1) is mis-specified, and equally efficient when (2.1) is correct i.e. $m(\cdot) = g_{\boldsymbol{\theta}_0}(\cdot)$. Further, it can also be shown that $\boldsymbol{\psi}_{\text{eff}}(\mathbf{Z})$ is the ‘efficient’ IF for estimating $\boldsymbol{\theta}_0$ under the semi-parametric model $\mathcal{M}_{\mathbf{X}} \equiv \{(\mathbb{P}_{Y|\mathbf{X}}, \mathbb{P}_{\mathbf{X}}) : \mathbb{P}_{\mathbf{X}} \text{ is known, } \mathbb{P}_{Y|\mathbf{X}} \text{ is unrestricted upto Assumption 2.1 (a)}\}$. Thus, $\hat{\boldsymbol{\theta}}_{np}$ also globally *achieves* the semi-parametric efficiency bound under $\mathcal{M}_{\mathbf{X}}$. Lastly, note that at any parametric sub-model in $\mathcal{M}_{\mathbf{X}}$ that corresponds to (2.1) being correct, $\hat{\boldsymbol{\theta}}$ also achieves optimality, thus showing that under $\mathcal{M}_{\mathbf{X}}$, it is not possible to improve upon $\hat{\boldsymbol{\theta}}$ if the linear model is correct.

REMARK 3.2. The asymptotic results in Theorem 3.1 require a kernel of order $q > p$ and h smaller in order than the ‘optimal’ bandwidth order $h_{opt} = O(n^{-1/(2q+p)})$. This *under-smoothing* requirement, often encountered in two-step estimators involving a first-step smoothing (Newey, Hsieh and Robins, 1998), generally results in sub-optimal performance of $\hat{m}(\cdot)$. The optimal under-smoothed bandwidth order for Theorem 3.1 is given by: $O(n^{-1/(q+p)})$.

3.2. SS Estimators Based on Semi-Non-Parametric (SNP) Imputation.

The simple and intuitive imputation strategy in Section 3.1 based on a fully non-parametric p -dimensional KS is however often undesirable in practice owing to the curse of dimensionality. In order to accommodate larger p , we now propose a more flexible SNP imputation method involving a dimension reduction, if needed, followed by a non-parametric calibration. An additional ‘refitting’ step is proposed to reduce the impact of bias from non-parametric estimation and possibly inadequate imputation due to dimension reduction. We also introduce some flexibility in terms of the smoothing methods, apart from KS, that can be used for the non-parametric calibration.

Let $r \leq p$ be a fixed positive integer and let $\mathbf{P}_r = [\mathbf{p}_1, \dots, \mathbf{p}_r]_{p \times r}$ be any rank r transformation matrix. Let $\mathbf{X}_{\mathbf{P}_r} = \mathbf{P}_r' \mathbf{X}$. Given (r, \mathbf{P}_r) , we may now consider approximating the regression function $\mathbb{E}(Y|\mathbf{X})$ by smoothing Y over the r dimensional $\mathbf{X}_{\mathbf{P}_r}$ instead of the original $\mathbf{X} \in \mathbb{R}^p$. In general, \mathbf{P}_r can be user-defined and data dependent. A few reasonable choices of \mathbf{P}_r are discussed in Section 5. If \mathbf{P}_r depends only on the distribution of \mathbf{X} , it may be assumed to be known given the SS setting considered. If \mathbf{P}_r also depends on the distribution of Y , then it needs to be estimated from \mathcal{L} and the smoothing needs to be performed using the estimated \mathbf{P}_r .

For approximating $\mathbb{E}(Y|\mathbf{X})$, we may consider *any* reasonable smoothing technique \mathcal{T} . Some examples of \mathcal{T} include KS, kernel machine regression and smoothing splines. Let $m(\mathbf{x}; \mathbf{P}_r)$ denote the ‘target function’ for smoothing Y over $\mathbf{X}_{\mathbf{P}_r}$ using \mathcal{T} . For notational simplicity, the dependence of $m(\mathbf{x}; \mathbf{P}_r)$ and other quantities on \mathcal{T} is suppressed throughout. For $\mathcal{T} := \text{KS}$, the appropriate target function is given by: $m(\mathbf{x}; \mathbf{P}_r) = m_{\mathbf{P}_r}(\mathbf{P}_r' \mathbf{x})$, where $m_{\mathbf{P}_r}(\mathbf{z}) \equiv \mathbb{E}(Y | \mathbf{X}_{\mathbf{P}_r} = \mathbf{z})$. For basis function expansion based methods, $m(\mathbf{x}; \mathbf{P}_r)$ will typically correspond to the L_2 projection of $m(\mathbf{x}) \equiv \mathbb{E}(Y | \mathbf{X} = \mathbf{x}) \in \mathcal{L}_2(\mathbb{P}_{\mathbf{X}})$ onto the functional space spanned by the basis functions associated with \mathcal{T} . The results in this section apply to any choice of \mathcal{T} that satisfies the required conditions. In Section 4, we provide more specific results for the implementation of our methods using $\mathcal{T} := \text{KS}$.

Note that we do *not* assume $m(\mathbf{x}; \mathbf{P}_r) = m(\mathbf{x})$ anywhere, and hence the name ‘semi-non-parametric’. Clearly, with $\mathbf{P}_r = I_p$ and $\mathcal{T} := \text{KS}$, it reduces

to a fully non-parametric approach. We next describe the two sub-steps involved in step (I) of the SNP imputation: (Ia) smoothing, and (Ib) refitting.

(Ia) Smoothing Step. With \mathbf{P}_r and $m(\mathbf{x}; \mathbf{P}_r)$ as defined above, let $\widehat{\mathbf{P}}_r$ and $\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$ respectively denote their estimators based on \mathcal{L} . In order to address potential overfitting issues in the subsequent steps, we further consider generalized versions of these estimators based on \mathbb{K} -fold CV for a given fixed integer $\mathbb{K} \geq 2$. For any $\mathbb{K} \geq 2$, let $\{\mathcal{L}_k\}_{k=1}^{\mathbb{K}}$ denote a random partition of \mathcal{L} into \mathbb{K} disjoint subsets of equal sizes, $n_{\mathbb{K}} = n/\mathbb{K}$, with index sets $\{\mathcal{I}_k\}_{k=1}^{\mathbb{K}}$. Let \mathcal{L}_k^- denote the set excluding \mathcal{L}_k with size $n_{\mathbb{K}}^- = n - n_{\mathbb{K}}$ and respective index set \mathcal{I}_k^- . Let $\widehat{\mathbf{P}}_{r,k}$ and $\widehat{m}_k(\mathbf{x}; \widehat{\mathbf{P}}_{r,k})$ denote the corresponding estimators based on \mathcal{L}_k^- . Further, for notational consistency, we define for $\mathbb{K} = 1$, $\mathcal{L}_k = \mathcal{L}_k^- = \mathcal{L}$; $\mathcal{I}_k = \mathcal{I}_k^- = \{1, \dots, n\}$; $n_{\mathbb{K}} = n_{\mathbb{K}}^- = n$; $\widehat{\mathbf{P}}_{r,k} = \widehat{\mathbf{P}}_r$ and $\widehat{m}_k(\mathbf{x}; \widehat{\mathbf{P}}_{r,k}) = \widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$.

(Ib) Refitting Step. In this step, we fit the linear model to \mathcal{L} using \mathbf{X} as predictors and the estimated $m(\mathbf{X}; \mathbf{P}_r)$ as an *offset*. To motivate this, we recall that the fully non-parametric imputation given in Section 3.1 consistently estimates $\mathbb{E}(Y|\mathbf{X})$, the L_2 projection onto a space that always contains the working model space, i.e. the linear span of $\vec{\mathbf{X}}$. This need not be true for the SNP imputation, since we do not assume $m(\mathbf{X}; \mathbf{P}_r) = m(\mathbf{X})$ necessarily. The refitting step essentially ‘adjusts’ for this so that the final imputation, combining the predictions from these two steps, targets a space that contains the working model space. In particular, for $\mathcal{T} := \text{KS}$ with $r < p$, this step is critical to remove potential bias due to inadequate imputation.

Interestingly, it turns out that the refitting step should *always* be performed, *even when* $m(\mathbf{X}; \mathbf{P}_r) = m(\mathbf{X})$. It plays a crucial role in reducing the bias of the resulting SS estimator due to the inherent bias from non-parametric curve estimation. In particular, for $\mathcal{T} := \text{KS}$ with *any* $r \leq p$, it ensures that a bandwidth of the optimal order can be used, thereby *eliminating the under-smoothing issue* as encountered in Section 3.1. The target parameter for the refitting step is simply the regression coefficient obtained from regressing the residual $Y - m(\mathbf{X}; \mathbf{P}_r)$ on \mathbf{X} and may be defined as: $\boldsymbol{\eta}_{\mathbf{P}_r}$, the solution in $\boldsymbol{\eta} \in \mathbb{R}^{(p+1)}$ to the equation: $\mathbb{E}[\vec{\mathbf{X}}\{Y - m(\mathbf{X}; \mathbf{P}_r) - \vec{\mathbf{X}}'\boldsymbol{\eta}\}] = \mathbf{0}$. For any $\mathbb{K} \geq 1$, we estimate $\boldsymbol{\eta}_{\mathbf{P}_r}$ as $\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}$, the solution in $\boldsymbol{\eta}$ to the equation:

$$(3.4) \quad n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \vec{\mathbf{X}}_i \{Y_i - \widehat{m}_k(\mathbf{X}_i; \widehat{\mathbf{P}}_{r,k}) - \vec{\mathbf{X}}_i' \boldsymbol{\eta}\} = \mathbf{0}.$$

For $\mathbf{X}_i \in \mathcal{L}_k$, the estimate of $m(\mathbf{X}_i; \mathbf{P}_r)$ to be used as an offset is obtained from $\widehat{m}_k(\cdot; \widehat{\mathbf{P}}_{r,k})$ that is based on data in \mathcal{L}_k^- . For $\mathbb{K} \geq 2$, with $\mathcal{L}_k^- \perp\!\!\!\perp \mathcal{L}_k$, the residuals are thus estimated in a cross-validated manner. For $\mathbb{K} = 1$ however,

$\widehat{m}(\cdot; \widehat{\mathbf{P}}_r)$ is estimated using the entire \mathcal{L} which can lead to considerable underestimation of the true residuals owing to over-fitting and consequently, substantial finite sample bias in the resulting SS estimator of $\boldsymbol{\theta}_0$. This bias can be effectively reduced by using the CV approach with $\mathbb{K} \geq 2$. We next estimate the *target function* for the SNP imputation given by:

$$(3.5) \quad \mu(\mathbf{x}; \mathbf{P}_r) = m(\mathbf{x}; \mathbf{P}_r) + \vec{\mathbf{x}}' \boldsymbol{\eta}_{\mathbf{P}_r} \text{ as:}$$

$$(3.6) \quad \widehat{\mu}(\mathbf{x}; \widehat{\mathcal{P}}_{r, \mathbb{K}}) = \mathbb{K}^{-1} \sum_{k=1}^{\mathbb{K}} \widehat{m}_k(\mathbf{x}; \widehat{\mathbf{P}}_{r, k}) + \vec{\mathbf{x}}' \widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})},$$

where $\widehat{\mathcal{P}}_{r, \mathbb{K}} = \{\widehat{\mathbf{P}}_{r, k}\}_{k=1}^{\mathbb{K}}$. For notational simplicity, we suppress throughout the inherent dependence of $\widehat{\mu}(\cdot; \cdot)$ itself on \mathbb{K} and $\{\mathcal{L}_k\}_{k=1}^{\mathbb{K}}$. Note that similar to $m(\mathbf{X}; \mathbf{P}_r)$, we also do *not* assume $\mu(\mathbf{X}; \mathbf{P}_r) = m(\mathbf{X})$. Apart from the geometric motivation for the refitting step and its technical role in bias reduction, it also generally ensures the condition: $\mathbb{E}[\vec{\mathbf{X}}\{Y - \mu(\mathbf{X}; \mathbf{P}_r)\}] = \mathbf{0}$, *regardless* of the true underlying $m(\mathbf{X})$. This condition is a key requirement for the asymptotic expansions, in Theorem 3.2, of our resulting SS estimators. Using $\widehat{\mu}(\cdot; \widehat{\mathcal{P}}_{r, \mathbb{K}})$, we now construct our final SS estimator as follows.

SS Estimator from SNP Imputation. In step (II), we fit the linear model to the SNP imputed unlabeled data: $\{[\widehat{\mu}(\mathbf{X}_j; \widehat{\mathcal{P}}_{r, \mathbb{K}}), \mathbf{X}'_j]'\}_{j=n+1, \dots, n+N}$ and obtain a SS estimator $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ of $\boldsymbol{\theta}_0$ given by:

$$(3.7) \quad \widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} \text{ is the solution in } \boldsymbol{\theta} \text{ to } \frac{1}{N} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \{\widehat{\mu}(\mathbf{X}_j; \widehat{\mathcal{P}}_{r, \mathbb{K}}) - \vec{\mathbf{X}}_j' \boldsymbol{\theta}\} = \mathbf{0}.$$

For convenience of further discussion, let us define: $\forall k \in \{1, \dots, \mathbb{K}\}$,

$$(3.8) \quad \widehat{\Delta}_k(\mathbf{x}; \mathbf{P}_r, \widehat{\mathbf{P}}_{r, k}) = \widehat{m}_k(\mathbf{x}; \widehat{\mathbf{P}}_{r, k}) - m(\mathbf{x}; \mathbf{P}_r) \quad \forall \mathbf{x} \in \mathcal{X}, \quad \text{and}$$

$$(3.9) \quad \widehat{\mathbf{G}}_k(\mathbf{x}) = \vec{\mathbf{x}} \widehat{\Delta}_k(\mathbf{x}; \mathbf{P}_r, \widehat{\mathbf{P}}_{r, k}) - \mathbb{E}_{\mathbf{X}}\{\vec{\mathbf{X}} \widehat{\Delta}_k(\mathbf{X}; \mathbf{P}_r, \widehat{\mathbf{P}}_{r, k})\} \quad \forall \mathbf{x} \in \mathcal{X},$$

where $\mathbb{E}_{\mathbf{X}}(\cdot)$ denotes expectation w.r.t. $\mathbf{X} \in \mathcal{U}$. The dependence of $\widehat{\mathbf{G}}_k(\cdot)$ on $(\mathbf{P}_r, \widehat{\mathbf{P}}_{r, k})$ and $\mathbb{P}_{\mathbf{X}}$ is suppressed here for notational simplicity. We now present our main result summarizing the properties of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$.

THEOREM 3.2. *Suppose that \mathcal{T} satisfies: (i) $\sup_{\mathbf{x} \in \mathcal{X}} |m(\mathbf{x}; \mathbf{P}_r)| < \infty$ and (ii) $\sup_{\mathbf{x} \in \mathcal{X}} |\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r) - m(\mathbf{x}; \mathbf{P}_r)| = O_p(c_n)$ for some $c_n = o(1)$. With $\widehat{\mathbf{G}}_k(\cdot)$ as in (3.9), define $\mathbb{G}_{n, \mathbb{K}} = n^{-\frac{1}{2}} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \widehat{\mathbf{G}}_k(\mathbf{X}_i)$. Then, for any $\mathbb{K} \geq 1$,*

$$(3.10) \quad n^{\frac{1}{2}} \left(\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0 \right) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}(\mathbf{Z}_i; \mathbf{P}_r) - \boldsymbol{\Gamma}^{-1} \mathbb{G}_{n, \mathbb{K}} + O_p(c_{n, \mathbb{K}}^*),$$

where $\psi(\mathbf{Z}; \mathbf{P}_r) = \mathbf{\Gamma}^{-1}[\vec{\mathbf{X}}\{Y - \mu(\mathbf{X}; \mathbf{P}_r)\}]$ and $c_{n, \mathbb{K}}^* = c_{n_{\mathbb{K}}} + n^{-\frac{1}{2}} + (n/N)^{\frac{1}{2}} = o(1)$. Further, for any fixed $\mathbb{K} \geq 2$, $\mathbb{G}_{n, \mathbb{K}} = O_p(c_{n_{\mathbb{K}}})$, so that

$$(3.11) \quad n^{\frac{1}{2}} \left(\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0 \right) = n^{-\frac{1}{2}} \sum_{i=1}^n \psi(\mathbf{Z}_i; \mathbf{P}_r) + O_p(c_{n_{\mathbb{K}}} + c_{n, \mathbb{K}}^*),$$

which converges in distribution to $\mathcal{N}_{(p+1)}[\mathbf{0}, \boldsymbol{\Sigma}\{\mu(\cdot; \mathbf{P}_r)\}]$.

REMARK 3.3. If the imputation is ‘sufficient’ so that $\mu(\mathbf{x}; \mathbf{P}_r) = m(\mathbf{x})$, then $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$, for any $\mathbb{K} \geq 2$, enjoys the same set of optimality properties as those noted in Remark 3.1 for $\widehat{\boldsymbol{\theta}}_{np}$ (while requiring less stringent assumptions about $K(\cdot)$ and h , if KS is used). If $\mu(\mathbf{x}; \mathbf{P}_r) \neq m(\mathbf{x})$, then it is however unclear whether $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ is always more efficient than $\widehat{\boldsymbol{\theta}}$. This will be addressed in Section 3.3 where we develop the final EASE.

REMARK 3.4. Apart from the fairly mild condition (i), Theorem 3.2 only requires uniform consistency of $\widehat{m}(\cdot; \widehat{\mathbf{P}}_r)$ w.r.t. $m(\cdot; \mathbf{P}_r)$ for establishing the $n^{\frac{1}{2}}$ -consistency and asymptotic normality (CAN) of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ for any $\mathbb{K} \geq 2$. The uniform consistency typically holds for a wide range of smoothing methods \mathcal{T} under fairly general conditions. For $\mathcal{T} := \text{KS}$ in particular, we provide explicit results in Section 4 under mild regularity conditions that allow the use of any kernel order and the associated optimal bandwidth order. This is a notable relaxation from the stringent requirements for Theorem 3.1 that necessitate under-smoothing and the use of higher order kernels.

REMARK 3.5. The CAN property of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, 1)}$ has *not* yet been established. The term $\mathbb{G}_{n, \mathbb{K}}$ in (3.10) behaves quite differently when $\mathbb{K} = 1$ compared to $\mathbb{K} \geq 2$. We derive the properties of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, 1)}$ in Section 4 for $\mathcal{T} := \text{KS}$.

3.3. *Efficient and Adaptive Semi-Supervised Estimators (EASE)*. To ensure adaptivity even when $\mu(\mathbf{x}; \mathbf{P}_r) \neq m(\mathbf{x})$, we now define the final EASE as an optimal linear combination of $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$. Specifically, for any fixed $(p+1) \times (p+1)$ matrix $\boldsymbol{\Delta}$, $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}(\boldsymbol{\Delta}) = \widehat{\boldsymbol{\theta}} + \boldsymbol{\Delta}(\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \widehat{\boldsymbol{\theta}})$ is a CAN estimator of $\boldsymbol{\theta}_0$ whenever $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ are, and an optimal $\boldsymbol{\Delta}$ can be selected easily to minimize the asymptotic variance of the combined estimator. For simplicity, we focus here on $\boldsymbol{\Delta}$ being a diagonal matrix with $\boldsymbol{\Delta} = \text{diag}(\delta_1, \dots, \delta_{p+1})$. Then the EASE is defined as $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E \equiv \widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}(\widehat{\boldsymbol{\Delta}})$ with $\widehat{\boldsymbol{\Delta}}$ being any consistent estimator (see Section 3.4 for details) of the

minimizer $\bar{\Delta} = \text{diag}(\bar{\delta}_1, \dots, \bar{\delta}_{p+1})$, where $\forall 1 \leq l \leq (p+1)$,

$$(3.12) \quad \bar{\delta}_l = - \lim_{\epsilon \downarrow 0} \frac{\text{Cov} \left\{ \psi_{0[l]}(\mathbf{Z}), \psi_{[l]}(\mathbf{Z}; \mathbf{P}_r) - \psi_{0[l]}(\mathbf{Z}) \right\}}{\text{Var} \left\{ \psi_{[l]}(\mathbf{Z}; \mathbf{P}_r) - \psi_{0[l]}(\mathbf{Z}) \right\} + \epsilon},$$

and for any vector \mathbf{a} , $\mathbf{a}_{[l]}$ denotes its l^{th} component. Note that in (3.12), the ϵ and the limit outside are included to formally account for the case: $\psi_{0[l]}(\mathbf{Z}) = \psi_{[l]}(\mathbf{Z}, \mathbf{P}_r)$ a.s. $[\mathbb{P}_{\mathbf{Z}}]$, when we define $\bar{\delta}_l = 0$ for identifiability.

It is straightforward to show that $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ and $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}(\bar{\Delta})$ are asymptotically equivalent, so that $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ is a RAL estimator of $\boldsymbol{\theta}_0$ satisfying:

$$n^{\frac{1}{2}} \left(\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E - \boldsymbol{\theta}_0 \right) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}(\mathbf{Z}_i; \mathbf{P}_r, \bar{\Delta}) + o_p(1) \xrightarrow{d} \mathcal{N}_{(p+1)}[\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{P}_r}(\bar{\Delta})],$$

as $n \rightarrow \infty$, where $\boldsymbol{\psi}(\mathbf{Z}; \mathbf{P}_r, \bar{\Delta}) = \boldsymbol{\psi}_0(\mathbf{Z}) + \bar{\Delta} \{ \boldsymbol{\psi}(\mathbf{Z}; \mathbf{P}_r) - \boldsymbol{\psi}_0(\mathbf{Z}) \}$ and $\boldsymbol{\Sigma}_{\mathbf{P}_r}(\bar{\Delta}) = \text{Var} \{ \boldsymbol{\psi}(\mathbf{Z}; \mathbf{P}_r, \bar{\Delta}) \}$. Note that when either the linear model holds or the SNP imputation is sufficient, then $\boldsymbol{\psi}(\mathbf{Z}; \mathbf{P}_r, \bar{\Delta}) = \boldsymbol{\psi}_{\text{eff}}(\mathbf{Z})$, so that $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ is asymptotically optimal in the sense of Remark 3.1. Further, when neither cases hold, $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ is no longer optimal, but is *still* efficient and adaptive compared to $\hat{\boldsymbol{\theta}}$. Lastly, if the imputation is certain to be sufficient (e.g. if $r = p$ and $\mathcal{T} := \text{KS}$), we may simply define $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E = \hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$.

REMARK 3.6. It can be shown that under $\mathcal{M}_{\mathbf{X}}$, defined in Remark 3.1, the class of all possible IFs achievable by RAL estimators of $\boldsymbol{\theta}_0$ is given by: $\mathcal{IF}_{\boldsymbol{\theta}_0, \mathcal{M}_{\mathbf{X}}} = \{ \boldsymbol{\phi}_g(\mathbf{Z}) \equiv \boldsymbol{\Gamma}^{-1} \bar{\mathbf{X}} \{ Y - g(\mathbf{X}) \} : g(\cdot) \in \mathcal{L}_2(\mathbb{P}_{\mathbf{X}}), \mathbb{E} \{ \boldsymbol{\phi}_g(\mathbf{Z}) \} = \mathbf{0} \}$. The IFs achieved by $\hat{\boldsymbol{\theta}}$, $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ and $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ are clearly members of this class. The SNP imputation, for various choices of the imputation function $\mu(\cdot; \mathbf{P}_r)$, therefore equips us with a *family* of RAL estimator pairs $\{ \hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}, \hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E \}$ for estimating $\boldsymbol{\theta}_0$. The IF of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ is further guaranteed to dominate that of $\hat{\boldsymbol{\theta}}$, and when $\mu(\cdot; \mathbf{P}_r) = m(\cdot)$, it also dominates all other IFs $\in \mathcal{IF}_{\boldsymbol{\theta}_0, \mathcal{M}_{\mathbf{X}}}$.

3.4. *Inference for EASE and the SNP Imputation Based SS Estimators.* We now provide procedures for making inference about $\boldsymbol{\theta}_0$ based on $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ and $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ obtained using $\mathbb{K} \geq 2$. We also employ a ‘double’ CV to overcome bias in variance estimation due to over-fitting. A key step involved in the variance estimation is to obtain reasonable estimates of $\{ \mu(\mathbf{X}_i; \mathbf{P}_r) \}_{i=1}^n$.

Although $\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}$ in (3.4) was constructed via CV, the corresponding estimate, $\widehat{\mu}(\mathbf{x}; \widehat{\mathcal{P}}_{r, \mathbb{K}})$ in (3.6), of $\mu(\mathbf{x}; \mathbf{P}_r)$ is likely to be over-fitted for $\mathbf{X}_i \in \mathcal{L}$. To construct bias corrected estimates of $\mu(\mathbf{X}_i; \mathbf{P}_r)$, we first obtain \mathbb{K} separate *doubly cross-validated* estimates of $\boldsymbol{\eta}_{\mathbf{P}_r}$, $\{\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}^k : k = 1, \dots, \mathbb{K}\}$, with $\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}^k$, for each k , being the solution in $\boldsymbol{\eta}$ to $\sum_{k' \neq k} \mathcal{S}_{k'}(\boldsymbol{\eta}) = \mathbf{0}$, where

$$\mathcal{S}_{k'}(\boldsymbol{\eta}) = \sum_{i \in \mathcal{I}_{k'}} \overrightarrow{\mathbf{X}}_i \{Y_i - \widehat{m}_{k'}(\mathbf{X}_i; \widehat{\mathbf{P}}_{r, k'}) - \overrightarrow{\mathbf{X}}_i' \boldsymbol{\eta}\} \quad \forall k' \in \{1, \dots, \mathbb{K}\}.$$

For each k and $k' \neq k$, $\mathcal{S}_{k'}(\boldsymbol{\eta})$ is constructed such that $\{\mathbf{Z}_i : i \in \mathcal{I}_{k'}\}$ used for obtaining $\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}^k$ is independent of $\widehat{m}_{k'}(\cdot; \widehat{\mathbf{P}}_{r, k'})$ that is based on $\mathcal{L}_{k'}^- \perp \mathcal{L}_{k'}$. Then, for each $\mathbf{X}_i \in \mathcal{L}_k$ and $k \in \{1, \dots, \mathbb{K}\}$, we may estimate $\mu(\mathbf{X}_i; \mathbf{P}_r)$ as:

$$\widehat{\mu}_k(\mathbf{X}_i; \widehat{\mathcal{P}}_{r, \mathbb{K}}) = \widehat{m}_k(\mathbf{X}_i; \widehat{\mathbf{P}}_{r, k}) + \overrightarrow{\mathbf{X}}_i' \widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}^k.$$

We exclude $\mathcal{S}_k(\boldsymbol{\eta})$ in the construction of $\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})}^k$ to reduce over-fitting bias in the residuals $\{Y_i - \widehat{\mu}_k(\mathbf{X}_i; \widehat{\mathcal{P}}_{r, \mathbb{K}})\}$ which we now use for estimating the IFs.

For each $\mathbf{Z}_i \in \mathcal{L}_k$ and $k \in \{1, \dots, \mathbb{K}\}$, we estimate $\boldsymbol{\psi}_0(\mathbf{Z}_i)$ and $\boldsymbol{\psi}(\mathbf{Z}_i; \mathbf{P}_r)$, the corresponding IFs of $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$, respectively as:

$$\widehat{\boldsymbol{\psi}}_0(\mathbf{Z}_i) = \widehat{\boldsymbol{\Gamma}}^{-1} \{\overrightarrow{\mathbf{X}}_i (Y_i - \overrightarrow{\mathbf{X}}_i' \widehat{\boldsymbol{\theta}})\} \text{ and } \widehat{\boldsymbol{\psi}}_k(\mathbf{Z}_i; \mathbf{P}_r) = \widehat{\boldsymbol{\Gamma}}^{-1} [\overrightarrow{\mathbf{X}}_i \{Y_i - \widehat{\mu}_k(\mathbf{X}_i; \widehat{\mathcal{P}}_{r, \mathbb{K}})\}],$$

where $\widehat{\boldsymbol{\Gamma}}$ denotes any consistent estimator of $\boldsymbol{\Gamma}$ from \mathcal{L} and/or \mathcal{U} , e.g. $\widehat{\boldsymbol{\Gamma}} = \boldsymbol{\Gamma}_n \equiv n^{-1} \sum_{i=1}^n \overrightarrow{\mathbf{X}}_i \overrightarrow{\mathbf{X}}_i'$ based on \mathcal{L} , or $\widehat{\boldsymbol{\Gamma}} = \boldsymbol{\Gamma}_N \equiv N^{-1} \sum_{j=n+1}^{n+N} \overrightarrow{\mathbf{X}}_j \overrightarrow{\mathbf{X}}_j'$ based on \mathcal{U} . Then, $\boldsymbol{\Sigma}\{\mu(\cdot; \mathbf{P}_r)\}$ in (3.11) may be consistently estimated as:

$$\widehat{\boldsymbol{\Sigma}}\{\mu(\cdot; \mathbf{P}_r)\} = n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \widehat{\boldsymbol{\psi}}_k(\mathbf{Z}_i; \mathbf{P}_r) \widehat{\boldsymbol{\psi}}_k'(\mathbf{Z}_i; \mathbf{P}_r).$$

To estimate the combination matrix $\overline{\boldsymbol{\Delta}}$ in (3.12) and the asymptotic variance, $\boldsymbol{\Sigma}_{\mathbf{P}_r}(\overline{\boldsymbol{\Delta}})$, of EASE consistently, let us define, $\forall 1 \leq l \leq (p+1)$,

$$\begin{aligned} \widehat{\sigma}_{l,12} &= -n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \widehat{\boldsymbol{\psi}}_{0[l]}(\mathbf{Z}_i) \{\widehat{\boldsymbol{\psi}}_{k[l]}(\mathbf{Z}_i; \mathbf{P}_r) - \widehat{\boldsymbol{\psi}}_{0[l]}(\mathbf{Z}_i)\}, \\ \widehat{\sigma}_{l,22} &= n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \{\widehat{\boldsymbol{\psi}}_{k[l]}(\mathbf{Z}_i; \mathbf{P}_r) - \widehat{\boldsymbol{\psi}}_{0[l]}(\mathbf{Z}_i)\}^2, \end{aligned}$$

and $\widehat{\delta}_l = \widehat{\sigma}_{l,12} / (\widehat{\sigma}_{l,22} + \epsilon_n)$ for some sequence $\epsilon_n \rightarrow 0$ with $n^{\frac{1}{2}} \epsilon_n \rightarrow \infty$. Then, we estimate $\overline{\boldsymbol{\Delta}}$ and $\boldsymbol{\Sigma}_{\mathbf{P}_r}(\overline{\boldsymbol{\Delta}})$ respectively as: $\widehat{\boldsymbol{\Delta}} = \text{diag}(\widehat{\delta}_1, \dots, \widehat{\delta}_{p+1})$ and

$$\widehat{\boldsymbol{\Sigma}}_{\mathbf{P}_r}(\widehat{\boldsymbol{\Delta}}) = n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \widehat{\boldsymbol{\psi}}_k(\mathbf{Z}_i; \mathbf{P}_r, \widehat{\boldsymbol{\Delta}}) \widehat{\boldsymbol{\psi}}_k'(\mathbf{Z}_i; \mathbf{P}_r, \widehat{\boldsymbol{\Delta}}),$$

where $\widehat{\boldsymbol{\psi}}_k(\mathbf{Z}; \mathbf{P}_r, \widehat{\boldsymbol{\Delta}}) = \widehat{\boldsymbol{\psi}}_0(\mathbf{Z}) + \widehat{\boldsymbol{\Delta}} \{\widehat{\boldsymbol{\psi}}_k(\mathbf{Z}; \mathbf{P}_r) - \widehat{\boldsymbol{\psi}}_0(\mathbf{Z})\} \quad \forall k \in \{1, \dots, \mathbb{K}\}$. Normal confidence intervals (CIs) for the parameters of interest can also be constructed accordingly based on these variance estimates.

4. Implementation Based on KS. We next detail the specific implementation of the SNP imputation based on KS estimators. With $\mathcal{T} := \text{KS}$, the target function for the smoothing is given by: $m(\mathbf{x}; \mathbf{P}_r) = m_{\mathbf{P}_r}(\mathbf{P}'_r \mathbf{x}) \equiv \mathbb{E}(Y | \mathbf{X}_{\mathbf{P}_r} = \mathbf{P}'_r \mathbf{x})$. For simplicity, we assume that $\mathbf{X}_{\mathbf{P}_r}$ is continuous with a density $f_{\mathbf{P}_r}(\cdot)$ and support $\mathcal{X}_{\mathbf{P}_r} \equiv \{\mathbf{P}'_r \mathbf{x} : \mathbf{x} \in \mathcal{X}\} \subseteq \mathbb{R}^r$. Let us now consider the following class of local constant KS estimators for $m(\mathbf{x}; \mathbf{P}_r)$:

$$(4.1) \quad \hat{m}_k(\mathbf{x}; \hat{\mathbf{P}}_{r,k}) = \frac{\frac{1}{n_{\mathbb{K}} h^r} \sum_{i \in \mathcal{I}_k^-} \{K_h(\hat{\mathbf{P}}'_{r,k} \mathbf{X}_i, \hat{\mathbf{P}}'_{r,k} \mathbf{x})\} Y_i}{\frac{1}{n_{\mathbb{K}} h^r} \sum_{i \in \mathcal{I}_k^-} K_h(\hat{\mathbf{P}}'_{r,k} \mathbf{X}_i, \hat{\mathbf{P}}'_{r,k} \mathbf{x})} \quad \forall 1 \leq k \leq \mathbb{K},$$

where $K_h(\cdot)$ and h are as in Section 3.1 with $K(\cdot)$ now being a suitable kernel on \mathbb{R}^r . In the light of Theorem 3.2, we focus primarily on establishing the uniform consistency of $\hat{m}(\mathbf{x}; \hat{\mathbf{P}}_r) \equiv \hat{m}_1(\mathbf{x}; \hat{\mathbf{P}}_{r,1})$ in (4.1) with $\mathbb{K} = 1$, *accounting* for the additional estimation error from $\hat{\mathbf{P}}_r$. For establishing the desired result, we shall assume the following regularity conditions to hold:

ASSUMPTION 4.1. (i) $K(\cdot)$ is a symmetric kernel of order $q \geq 2$ with finite q^{th} moments. (ii) $K(\cdot)$ is bounded, integrable and is either Lipschitz continuous with a compact support or, has a bounded derivative $\nabla K(\cdot)$ which satisfies: $\|\nabla K(\mathbf{z})\| \leq \Lambda \|\mathbf{z}\|^{-\rho} \forall \mathbf{z} \in \mathbb{R}^r$ with $\|\mathbf{z}\| > L$, where $\Lambda > 0$, $L > 0$ and $\rho > 1$ are some fixed constants, and $\|\cdot\|$ denotes the standard L_2 vector norm. (iii) $\mathcal{X}_{\mathbf{P}_r} \subseteq \mathbb{R}^r$ is compact. $\mathbb{E}(|Y|^s) < \infty$ for some $s > 2$. $\mathbb{E}(|Y|^s | \mathbf{X}_{\mathbf{P}_r} = \mathbf{z}) f_{\mathbf{P}_r}(\mathbf{z})$ and $f_{\mathbf{P}_r}(\mathbf{z})$ are bounded on $\mathcal{X}_{\mathbf{P}_r}$. (iv) $f_{\mathbf{P}_r}(\mathbf{z})$ is bounded away from 0 on $\mathcal{X}_{\mathbf{P}_r}$. (v) $m_{\mathbf{P}_r}(\mathbf{z})$ and $f_{\mathbf{P}_r}(\mathbf{z})$ are both q times continuously differentiable with bounded q^{th} derivatives on some open set $\mathcal{X}_{0, \mathbf{P}_r} \supseteq \mathcal{X}_{\mathbf{P}_r}$. *Additional Conditions (required only when \mathbf{P}_r needs to be estimated)*: (vi) $K(\cdot)$ has a bounded and integrable derivative $\nabla K(\cdot)$. (vii) $\nabla K(\cdot)$ satisfies: $\|\nabla K(\mathbf{z}_1) - \nabla K(\mathbf{z}_2)\| \leq \|\mathbf{z}_1 - \mathbf{z}_2\| \phi(\mathbf{z}_1) \forall \mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^r$ such that $\|\mathbf{z}_1 - \mathbf{z}_2\| \leq L^*$, for some fixed constant $L^* > 0$, and some bounded and integrable function $\phi : \mathbb{R}^r \rightarrow \mathbb{R}^+$. (viii) $\nabla K(\cdot)$ is Lipschitz continuous on \mathbb{R}^r . (ix) $\mathbb{E}(\mathbf{X} | \mathbf{X}_{\mathbf{P}_r} = \mathbf{z})$ and $\mathbb{E}(\mathbf{X}Y | \mathbf{X}_{\mathbf{P}_r} = \mathbf{z})$ are both continuously differentiable with bounded first derivatives on $\mathcal{X}_{0, \mathbf{P}_r} \supseteq \mathcal{X}_{\mathbf{P}_r}$.

Assumption 4.1, mostly adopted from Hansen (2008), imposes some mild smoothness and moment conditions most of which are fairly standard, except perhaps the conditions on $K(\cdot)$ in (vi)-(viii) all of which are however satisfied by the Gaussian kernel among others. We now propose the following result.

THEOREM 4.1. *Suppose $(\hat{\mathbf{P}}_r - \mathbf{P}_r) = O_p(\alpha_n)$ for some $\alpha_n = o(1)$ with $\alpha_n = 0$ identically if \mathbf{P}_r is known. Let q be the order of the kernel $K(\cdot)$ in*

(4.1) for some integer $q \geq 2$. Define:

$$a_{n,1} = \alpha_n \left(\frac{\log n}{nh^{r+2}} \right)^{\frac{1}{2}} + \alpha_n^2 h^{-(r+2)} + \alpha_n, \quad a_{n,2} = \left(\frac{\log n}{nh^r} \right)^{\frac{1}{2}} + h^q$$

and assume that each of the terms involved in $a_{n,1} = o(1)$ and $a_{n,2} = o(1)$. Then, under Assumption 4.1, $\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$, based on (4.1), satisfies:

$$(4.2) \quad \sup_{\mathbf{x} \in \mathcal{X}} |\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r) - m(\mathbf{x}; \mathbf{P}_r)| = O_p(a_{n,1} + a_{n,2}).$$

REMARK 4.1. Theorem 4.1 establishes the L_∞ error rate of $\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$ under mild regularity conditions and restrictions on h . Among its various implications, the rate also ensures uniform consistency of $\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$ at the optimal bandwidth order: $h_{opt} = O(n^{-1/(2q+r)})$ for any kernel order $q \geq 2$ and any $r \leq p$, as long as $\alpha_n = o(n^{-(r+2)/(4q+2r)})$ which always includes: $\alpha_n = O(n^{-\frac{1}{2}})$ and $\alpha_n = 0$. These two cases are particularly relevant in practice as \mathbf{P}_r being finite dimensional, $n^{\frac{1}{2}}$ -consistent estimators of \mathbf{P}_r should typically exist. For both cases, using h_{opt} results in $a_{n,1}$ to be of lower order (for $q > 2$) or the same order (for $q = 2$) compared to that of the main term $a_{n,2}$, so that the usual optimal rate prevails as the overall error rate.

Properties of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ for $\mathbb{K} = 1$. We now address the CAN property of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ for $\mathbb{K} = 1$ under the KS framework. Based on (3.10) and Remark 3.5, the only step required for this is to effectively control the term $\mathbb{G}_{n, \mathbb{K}}$ in (3.10). We propose the following result in this regard.

THEOREM 4.2. *Let $\mathbb{K} = 1$, $\mathcal{T} := KS$, $\mathbb{G}_{n, \mathbb{K}}$ be as in (3.10), and $\widehat{m}(\mathbf{x}; \widehat{\mathbf{P}}_r)$ be the KS estimator based on (4.1). Let α_n , $a_{n,1}$ and $a_{n,2}$ be as in Theorem 4.1 with $(\widehat{\mathbf{P}}_r - \mathbf{P}_r) = O_p(\alpha_n)$. Assume that $a_{n,1}^*$ and $a_{n,2}^*$ are $o(1)$, where*

$$a_{n,1}^* = \alpha_n + \frac{\alpha_n}{n^{\frac{1}{2}} h^{(r+1)}} + n^{\frac{1}{2}} \alpha_n^2 h^{-2} + n^{\frac{1}{2}} a_{n,1}^2 + n^{\frac{1}{2}} a_{n,1} a_{n,2} \quad \text{and} \quad a_{n,2}^* = n^{\frac{1}{2}} a_{n,2}^2.$$

Then, under Assumption 4.1, $\mathbb{G}_{n, \mathbb{K}} = O_p(a_{n,1}^* + a_{n,2}^*) = o_p(1)$. Further, let $c_{n, \mathbb{K}}^*$ be as in Theorem 3.2 with $c_n = (a_{n,1} + a_{n,2})$. Then, using (3.10),

$$(4.3) \quad n^{\frac{1}{2}} \left(\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0 \right) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}(\mathbf{Z}_i, \mathbf{P}_r) + O_p(c_{n, \mathbb{K}}^* + d_n),$$

where $d_n = a_{n,1}^* + a_{n,2}^*$. Hence, $n^{\frac{1}{2}} (\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0) \xrightarrow{d} \mathcal{N}_{(p+1)}[\mathbf{0}, \boldsymbol{\Sigma}\{\mu(\cdot; \mathbf{P}_r)\}]$.

REMARK 4.2. Note that the term $a_{n,2}^*$ *always* requires $q > r/2$ in order to converge to 0, thus showing the contrasting behavior of the case $\mathbb{K} = 1$ compared to $\mathbb{K} \geq 2$ where no such higher order kernel restriction is required. Nevertheless, when $\alpha_n = O(n^{-\frac{1}{2}})$ or $\alpha_n = 0$, the optimal bandwidth order: $h_{opt} = O(n^{-1/(2q+r)})$ can indeed be *still* used as long as $q > r/2$ is satisfied. Despite these facts and all the theoretical guarantee in Theorem 4.2, empirical evidence however seems to suggest that $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r,1)}$ can be substantially *biased* in finite samples, in part due to over-fitting.

REMARK 4.3. *Technical benefits of refitting and CV:* Suppose that $\mathbf{P}_r = I_p$, so that the SNP imputation with $\mathcal{T} := \text{KS}$ is indeed sufficient. Further, assume that all of Theorems 3.1-4.2 hold, so that the estimators $\hat{\boldsymbol{\theta}}_{np}$, $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r,1)}$, and $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r,\mathbb{K})}$ ($\mathbb{K} \geq 2$) are comparable and all asymptotically optimal. However, their constructions are quite different which can significantly affect their finite sample performances. $\hat{\boldsymbol{\theta}}_{np}$ is based on KS only, and requires stringent under-smoothing and a kernel of order $q > p$ (Remark 3.2); $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r,1)}$ is based on KS and refitting (*although* the KS itself is certain to be sufficient), and requires no under-smoothing but needs a (weaker) kernel order condition ($q > p/2$) (Remark 4.2); while $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r,\mathbb{K})}$ ($\mathbb{K} \geq 2$) additionally involves CV, and requires no under-smoothing or higher order kernel conditions (Remark 3.4). This highlights the critical role played by refitting and CV, apart from their primary roles in the SNP imputation, in removing any under-smoothing and/or higher order kernel restrictions when $\mathcal{T} := \text{KS}$, and this continues to hold for any other (r, \mathbf{P}_r) as well. In particular, it shows, rather surprisingly, that refitting should be performed in order to avoid under-smoothing *even if* the smoothing is known to be sufficient.

REMARK 4.4. As mentioned in Section 3.2, $\mathcal{T} := \text{KS}$ along with possible dimension reduction is just *one* reasonable choice of \mathcal{T} for implementing the SNP imputation, all technical requirements for which have been thoroughly established in Section 4. In general, other smoothing methods, as long as the requirements are satisfied, can also be equally used as choices of \mathcal{T} . One such choice could be kernel machine (KM) regression (with possibly no use of dimension reduction, as KM uses penalization to effectively regularize the target even with $\mathbf{P}_r = I_p$). We leave its implementation details to the reader as they are readily available in a multitude of references, and also skip any theoretical treatment, considering the primary goal and scope of this paper. However, detailed numerical results are presented in Section 6 for this choice of \mathcal{T} as well to illustrate the wider applicability of our proposed methods.

5. Choices of \mathbf{P}_r : Dimension Reduction Techniques. We next discuss choosing and estimating the matrix \mathbf{P}_r ($r < p$) to be used for dimension reduction, if required, in the SNP imputation, and which can play an important role in the sufficiency of the imputation. Simple choices of \mathbf{P}_r include r leading principal component directions of \mathbf{X} or any r canonical directions of \mathbf{X} . Note that under the SS setting, \mathbf{P}_r is effectively known if it only involves the distribution of \mathbf{X} , as is true for these choices. We now focus primarily on the case where \mathbf{P}_r also depends on the distribution of Y and hence, is unknown. Such a choice of \mathbf{P}_r is often desirable to ensure that the imputation is as ‘sufficient’ as possible for predicting Y . Several reasonable choices of such \mathbf{P}_r and their estimation are possible based on sufficient dimension reduction (s.d.r.) methods like Sliced Inverse Regression (SIR) (Li, 1991), Principal Hessian Directions (PHD) (Li, 1992; Cook, 1998), Sliced Average Variance Estimation (SAVE) (Cook and Weisberg, 1991; Cook and Lee, 1999) etc.

In particular, we focus here on SIR where the choice of \mathbf{P}_r is given by: $\mathbf{P}_r^0 = \Sigma^{-\frac{1}{2}} \bar{\mathbf{P}}_r$, with $\bar{\mathbf{P}}_r$ being the r leading eigenvectors of $\mathbb{M} = \text{Var}\{\mathbb{E}(\mathbb{X} | Y)\}$, where $\mathbb{X} = \Sigma^{-\frac{1}{2}}(\mathbf{X} - \boldsymbol{\mu})$, with $\boldsymbol{\mu} = \mathbb{E}(\mathbf{X})$, denotes the standardized version of \mathbf{X} . It is well known (Li, 1991) that these directions lead to an optimal (in some appropriate sense) r -dimensional linear transformation of \mathbf{X} that can be predicted by Y . Apart from these general optimality, they also have deeper implications in the context of s.d.r. We refer the reader to Li (1991) and other relevant references in the s.d.r. literature for further details.

For estimating \mathbf{P}_r^0 , we consider the SIR algorithm of Li (1991) and further propose a SS modification to it. With \mathbb{K} and $\{\mathcal{L}_k^-, \mathcal{I}_k^-, \hat{\mathbf{P}}_{r,k}\}_{k=1}^{\mathbb{K}}$ as before, let $(\hat{\boldsymbol{\mu}}_k, \hat{\boldsymbol{\Sigma}}_k)$ denote the estimates of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ based on \mathcal{L}_k^- and define $\mathbb{X}^{(k)} = \hat{\boldsymbol{\Sigma}}_k^{-\frac{1}{2}}(\mathbf{X} - \hat{\boldsymbol{\mu}}_k)$. Then, the original SIR algorithm estimates \mathbf{P}_r^0 based on \mathcal{L}_k^- as follows: (i) Divide the range of $\{Y_i\}_{i \in \mathcal{I}_k^-}$ into H slices $\{I_1, \dots, I_H\}$, where H may depend on $n_{\mathbb{K}}^-$. For $1 \leq h \leq H$, let $\hat{p}_{h,k}$ denote the proportion of $\{Y_i\}_{i \in \mathcal{I}_k^-}$ in slice I_h ; (ii) For each I_h , let $\hat{\mathbb{M}}_{h,k}$ denote the sample average of the set: $\{\mathbb{X}_i^{(k)} \in \mathcal{L}_k^- : Y_i \in I_h\}$; (iii) Estimate \mathbb{M} as: $\hat{\mathbb{M}}_k = \sum_{h=1}^H \hat{p}_{h,k} \hat{\mathbb{M}}_{h,k} \hat{\mathbb{M}}_{h,k}'$ and \mathbf{P}_r^0 as: $\hat{\mathbf{P}}_{r,k}^0 = \hat{\boldsymbol{\Sigma}}_k^{-\frac{1}{2}} \hat{\mathbf{P}}_{r,k}$, where $\hat{\mathbf{P}}_{r,k}$ denotes the r leading eigenvectors of $\hat{\mathbb{M}}_k$. However, the SIR algorithm often tends to give unstable estimates of \mathbf{P}_r^0 , especially for the directions corresponding to the smaller eigenvalues of \mathbb{M} . To improve the efficiency in estimating \mathbf{P}_r^0 , we now propose a semi-supervised SIR (SS-SIR) algorithm as follows.

SS-SIR Algorithm. Given $\{\mathcal{L}_k^-, \mathcal{I}_k^-, \hat{\mathbf{P}}_{r,k}\}_{k=1}^{\mathbb{K}}$, let $(\hat{\boldsymbol{\mu}}_k^*, \hat{\boldsymbol{\Sigma}}_k^*)$ denote the estimates of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ based on $\mathcal{L}_k^- \cup \mathcal{U}$ and define $\mathbb{X}^{(k^*)} = \hat{\boldsymbol{\Sigma}}_k^{*- \frac{1}{2}}(\mathbf{X} - \hat{\boldsymbol{\mu}}_k^*)$. Then

the SS-SIR proceeds as follows. Step (i) stays the same as in SIR. In step (ii), for each k , and each $j \in \{n+1, \dots, n+N\}$, we impute Y_j as $Y_{j,k}^* = Y_{\hat{i}_{j,k}}$, where $\hat{i}_{j,k} = \operatorname{argmin}_{i \in \mathcal{I}_k^-} \|\mathbb{X}_i^{(k^*)} - \mathbb{X}_j^{(k^*)}\|^2$. For each I_h , let $\widehat{\mathbb{M}}_{h,k}^*$ be the sample average of the set: $\{\mathbb{X}_i^{(k^*)} \in \mathcal{L}_k^- : Y_i \in I_h\} \cup \{\mathbb{X}_j^{(k^*)} \in \mathcal{U} : Y_{j,k}^* \in I_h\}$. Then in step (iii), we estimate \mathbb{M} as: $\widehat{\mathbb{M}}_k^* = \sum_{h=1}^H \widehat{p}_{h,k} \widehat{\mathbb{M}}_{h,k}^* \widehat{\mathbb{M}}_{h,k}^{*'}$ and then, \mathbf{P}_r^0 as: $\widehat{\mathbf{P}}_{r,k}^{0*} = \widehat{\Sigma}_k^{*-1/2} \widehat{\mathbf{P}}_{r,k}^*$, where $\widehat{\mathbf{P}}_{r,k}^*$ denotes the r leading eigenvectors of $\widehat{\mathbb{M}}_k^*$.

The SS-SIR algorithm aims to improve the estimation of \mathbf{P}_r^0 by making use of \mathcal{U} in step (ii) through a nearest neighbour approximation for the unobserved Y in \mathcal{U} using \mathcal{L}_k^- . With $n_{\mathbb{K}}^-$ large enough and $m(\cdot)$ smooth enough, the imputed and the true underlying Y should belong to the same slice with a high probability. Thus, the set of \mathbb{X} 's belonging to a particular slice is now 'enriched' and consequently, improved estimation of \mathbb{M} and \mathbf{P}_r^0 is expected. The proposed method based on a nearest neighbor approximation is also highly scalable and while other smoothing based approximations may be used, they can be computationally intensive. The SS-SIR algorithm is fairly robust to the choice of H , and $H = O(n^{1/2} \log n)$ seems to give fairly satisfactory performance. The slices may be chosen to have equal width or equal number of observations. For SIR, $n^{1/2}$ -consistency of the estimates are well established (Li, 1991; Duan and Li, 1991; Zhu and Ng, 1995) for various formulations under fairly general settings (without any model based assumptions). The theoretical properties of SS-SIR, although not derived here, are expected to follow similarly. Our simulation results (not shown here) further suggest that SS-SIR significantly outperforms SIR, leading to substantially improved estimation of $\boldsymbol{\theta}_0$ from the proposed methods.

6. Numerical Studies.

6.1. *Simulation Studies.* We conducted extensive simulation studies to examine the finite sample performance of our proposed point and interval estimation procedures as well as to compare with existing methods. Throughout we let $n = 500$, $N = 10000$, and considered $p = 2, 10$ and 20 . For our CV based methods, we let $\mathbb{K} = 5$. The true values of the target parameter $\boldsymbol{\theta}_0$ were estimated via monte carlo with a large sample size of $50,000$. For each configuration, the results were summarized based on 500 replications. Results for $p = 2$ are summarized in the [Supplementary Material](#), and the discussions below focus primarily on $p = 10$ and 20 .

We generated $\mathbf{X} \sim \mathcal{N}_p[\mathbf{0}, I_p]$ and restricted \mathbf{X} to $[-5, 5]^p$ to ensure its boundedness. Given $\mathbf{X} = \mathbf{x}$, we generated $Y \sim \mathcal{N}_1[m(\mathbf{x}), 1]$, where we considered four different choices of $m(\mathbf{x})$:

- (i) *Linear*: $m(\mathbf{x}) = \mathbf{x}'\mathbf{b}_p$;
- (ii) *Non-linear one component (NL1C)*: $m(\mathbf{x}) = (\mathbf{x}'\mathbf{b}_p) + (\mathbf{x}'\mathbf{b}_p)^2$;
- (iii) *Non-linear two component (NL2C)*: $m(\mathbf{x}) = (\mathbf{x}'\mathbf{b}_p)(1 + \mathbf{x}'\boldsymbol{\delta}_p)$; and
- (iv) *Non-linear three component (NL3C)*: $m(\mathbf{x}) = (\mathbf{x}'\mathbf{b}_p)(1 + \mathbf{x}'\boldsymbol{\delta}_p) + (\mathbf{x}'\boldsymbol{\omega}_p)^2$;

where, for each setting, we considered $\mathbf{b}_p = \mathbf{b}_p^{(1)} \equiv (\mathbf{1}'_{p/2}, \mathbf{0}'_{p/2})'$ and $\mathbf{b}_p = \mathbf{b}_p^{(2)} \equiv \mathbf{1}_p$, and set $\boldsymbol{\delta}_p = (\mathbf{0}'_{p/2}, \mathbf{1}'_{p/2})'$ and $\boldsymbol{\omega}_p = (1, 0, 1, 0, \dots, 1, 0)'_{p \times 1}$, where for any a , $\mathbf{1}_a = (1, \dots, 1)'_{a \times 1}$ and $\mathbf{0}_a = (0, \dots, 0)'_{a \times 1}$. Through appropriate choices of \mathbf{b}_p , $\boldsymbol{\delta}_p$ and $\boldsymbol{\omega}_p$, as applicable, these models can incorporate commonly encountered linear, quadratic and interaction effects.

For each setting, we used two choices of the smoothing method: (a) $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$ denoting KS with 2-dimensional smoothing over $\mathbf{P}'_r \mathbf{X} \equiv \mathbf{P}'_2 \mathbf{X}$, where \mathbf{P}_2 was estimated via SIR with $H = 100$ slices of equal width, following which $\{\hat{m}_k(\mathbf{x}, \hat{\mathbf{P}}_{r,k})\}_{k=1}^{\mathbb{K}}$ were obtained via KS using a Gaussian kernel; (b) $\mathcal{T} := \text{KM}$ where we let $\mathbf{P}_r = I_p$ and then estimated $\{\hat{m}_k(\mathbf{x}; I_p)\}_{k=1}^{\mathbb{K}}$ using kernel machine (KM) regression based on a radial basis function (RBF) kernel. Throughout, h for KS, and all tuning parameters for KM were selected via least squares CV. For (a) with $\mathbf{X} \sim \mathcal{N}_p[\mathbf{0}, I_p]$, results from Li (1991) imply that the SNP imputation with $r = 2$ is sufficient for models (i)-(iii), and insufficient for model (iv). For comparison, we also implemented two other SS estimators: the density ratio based ‘‘DRESS’’ estimator of Kawakita and Kanamori (2013) and the estimator of Sokolovska, Cappé and Yvon (2008) called ‘‘MSSL’’ by Kawakita and Kanamori (2013). The density ratio estimation for the DRESS estimator was implemented using either (i) linear bases $\{1, (\mathbf{X}_{[j]})^p_{j=1}\}$ (DRESS₁); or (ii) cubic bases $\{1, (\mathbf{X}_{[j]}^d)_{j=1, d=1}^{p, 3}\}$ (DRESS₃).

First, we compare the various estimators with respect to their efficiencies based on empirical mean squared error. In Table 1, we present the efficiencies of the proposed SNP and EASE estimators as well as other SS estimators relative to the OLS. As expected, under model mis-specification, our estimators are substantially more efficient than the OLS with the relative efficiency (RE) as high as near 5 fold when $p = 10$ and 3 fold when $p = 20$, for the non-linear models. The efficiency gain is generally lower for $p = 20$ than for $p = 10$, likely a consequence of overfitting of the non-parametric estimators involved in the SNP imputation for larger p . Comparing EASE to SNP, the EASE generally perform better for both linear and non-linear settings, as expected. Comparing the two smoothers, it appears that $\mathcal{T} := \text{KM}$ generally attains higher efficiency compared to that of $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$. This is in part due to the high variability in the SIR direction estimation which impacts performance of the resulting SS estimator in finite samples. Interestingly, none of the existing SS estimators perform well with REs ranging only from

about 0.9 to 1.1 across all settings.

(a) $p = 10$

Setting	Models	OLS (Ref.)	SNP EASE		SNP EASE		Other SS Estimators		
			$(\mathcal{T} := \text{KS}_{2, \mathbf{P}_2})$		$(\mathcal{T} := \text{KM})$		DRESS ₁	DRESS ₃	MSSL
(I)	Linear	1	0.895	0.983	0.772	0.985	0.982	0.927	0.982
	NL1C	1	4.481	4.424	4.501	5.543	1.136	1.110	1.135
	NL2C	1	2.683	2.700	4.268	5.055	1.120	1.016	1.119
	NL3C	1	2.772	2.795	4.481	5.560	1.102	1.025	1.103
(II)	Linear	1	0.841	0.989	0.657	0.993	0.981	0.924	0.981
	NL1C	1	4.511	4.585	4.416	5.471	1.132	1.030	1.130
	NL2C	1	3.596	3.634	4.405	5.497	1.127	1.042	1.128
	NL3C	1	3.280	3.301	4.636	5.566	1.110	1.079	1.109

(b) $p = 20$

Setting	Models	OLS (Ref.)	SNP EASE		SNP EASE		Other SS Estimators		
			$(\mathcal{T} := \text{KS}_{2, \mathbf{P}_2})$		$(\mathcal{T} := \text{KM})$		DRESS ₁	DRESS ₃	MSSL
(I)	Linear	1	0.673	0.986	0.740	0.981	0.956	0.866	0.956
	NL1C	1	2.256	2.288	2.680	3.630	1.035	0.920	1.035
	NL2C	1	1.414	1.388	2.661	3.544	1.032	0.922	1.033
	NL3C	1	1.539	1.531	2.605	3.510	1.049	0.931	1.051
(II)	Linear	1	0.519	0.991	0.609	0.989	0.958	0.872	0.958
	NL1C	1	2.290	2.346	2.669	3.660	1.032	0.908	1.031
	NL2C	1	1.899	1.917	2.766	3.963	1.036	0.917	1.036
	NL3C	1	1.937	1.949	2.682	3.702	1.046	0.958	1.046

Table 1: Efficiencies of $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ (SNP) and $\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ (EASE) using $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$ or $\mathcal{T} := \text{KM}$, as well as DRESS₁, DRESS₃ and MSSL, relative to $\widehat{\boldsymbol{\theta}}$ (OLS) with respect to the empirical mean squared error (MSE) under models (i), (ii), (iii) and (iv) each with: (I) $\mathbf{b}_p = \mathbf{b}_p^{(1)}$ or, (II) $\mathbf{b}_p = \mathbf{b}_p^{(2)}$.

We next examine the performance of the proposed inference procedures. In Table 2(a) and (b), we present the bias, empirical standard error (ESE), the average of the estimated standard error (ASE) and the coverage probability (CovP) of the 95% CIs for each component of $\boldsymbol{\theta}_0$ when $p = 10$ under the linear and NL2C models. In general, the EASE with both the KS and the KM smoothers have negligible biases although the KM based estimator appears to have slightly lower biases. The ASEs are close to the ESEs and the CovPs are close to the nominal level, suggesting that the variance estimators work well in practice with $\mathbb{K} = 5$. As shown in Table 2(c), the other SS estimators tend to have slightly larger biases and substantially larger standard errors (SEs) compared to our estimators under the NL2C model.

(a) OLS and EASE for the linear model.

Parameter	OLS ($\hat{\theta}$)		EASE ($\hat{\theta}_{(\mathbf{P}_r, \mathbb{K})}^E; \mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$)				EASE ($\hat{\theta}_{(\mathbf{P}_r, \mathbb{K})}^E; \mathcal{T} := \text{KM}$)			
	Bias	ESE	Bias	ESE	ASE	CovP	Bias	ESE	ASE	CovP
$\alpha_0 = 0$	-0.001	0.043	-0.001	0.043	0.044	0.95	0.000	0.043	0.044	0.96
$\beta_{01} = 1$	0.002	0.044	-0.003	0.045	0.044	0.94	0.004	0.047	0.044	0.93
$\beta_{02} = 1$	-0.001	0.044	-0.005	0.044	0.044	0.94	0.000	0.045	0.044	0.95
$\beta_{03} = 1$	-0.001	0.046	-0.005	0.046	0.044	0.95	-0.004	0.045	0.044	0.94
$\beta_{04} = 1$	-0.002	0.045	-0.006	0.045	0.044	0.94	0.001	0.047	0.044	0.94
$\beta_{05} = 1$	-0.004	0.048	-0.008	0.049	0.044	0.92	-0.001	0.046	0.044	0.95
$\beta_{06} = 0$	-0.000	0.045	-0.001	0.045	0.044	0.94	0.001	0.045	0.044	0.95
$\beta_{07} = 0$	0.003	0.046	0.003	0.046	0.044	0.93	0.001	0.043	0.044	0.96
$\beta_{08} = 0$	-0.001	0.045	-0.001	0.045	0.044	0.95	-0.000	0.048	0.044	0.94
$\beta_{09} = 0$	-0.002	0.047	-0.002	0.048	0.044	0.94	0.000	0.045	0.045	0.95
$\beta_{010} = 0$	0.003	0.045	0.003	0.045	0.044	0.94	-0.002	0.045	0.045	0.94

(b) OLS and EASE for the NL2C model.

Parameter	OLS ($\hat{\theta}$)		EASE ($\hat{\theta}_{(\mathbf{P}_r, \mathbb{K})}^E; \mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$)				EASE ($\hat{\theta}_{(\mathbf{P}_r, \mathbb{K})}^E; \mathcal{T} := \text{KM}$)			
	Bias	ESE	Bias	ESE	ASE	CovP	Bias	ESE	ASE	CovP
$\alpha_0 = 0$	-0.015	0.239	-0.016	0.146	0.136	0.93	0.013	0.105	0.096	0.93
$\beta_{01} = 1$	0.000	0.260	0.015	0.159	0.160	0.96	-0.004	0.124	0.112	0.93
$\beta_{02} = 1$	-0.004	0.269	0.017	0.173	0.158	0.93	0.010	0.127	0.113	0.93
$\beta_{03} = 1$	-0.015	0.249	0.018	0.156	0.158	0.95	-0.000	0.118	0.113	0.95
$\beta_{04} = 1$	-0.001	0.267	0.016	0.164	0.159	0.94	0.007	0.124	0.113	0.93
$\beta_{05} = 1$	0.013	0.260	0.019	0.164	0.158	0.94	0.002	0.120	0.113	0.94
$\beta_{06} = 0$	-0.010	0.281	0.008	0.164	0.155	0.94	0.005	0.119	0.112	0.94
$\beta_{07} = 0$	0.006	0.277	0.002	0.166	0.155	0.93	0.011	0.116	0.111	0.95
$\beta_{08} = 0$	-0.008	0.277	-0.004	0.167	0.156	0.94	-0.001	0.120	0.112	0.95
$\beta_{09} = 0$	0.002	0.279	0.003	0.160	0.157	0.95	0.007	0.118	0.113	0.95
$\beta_{010} = 0$	-0.008	0.272	0.002	0.160	0.155	0.95	0.004	0.130	0.111	0.91

(c) All other SS estimators for the models in (a) and (b) above.

Linear Model						NL2C Model					
DRESS ₁		DRESS ₃		MSSL		DRESS ₁		DRESS ₃		MSSL	
Bias	ESE	Bias	ESE	Bias	ESE	Bias	ESE	Bias	ESE	Bias	ESE
-0.001	0.043	-0.001	0.044	-0.001	0.043	-0.004	0.223	-0.003	0.226	-0.004	0.223
-0.002	0.044	-0.001	0.046	-0.002	0.044	-0.014	0.266	-0.009	0.279	-0.014	0.266
0.000	0.045	0.001	0.046	0.000	0.045	0.005	0.257	0.006	0.266	0.006	0.257
0.006	0.045	0.006	0.047	0.006	0.045	-0.013	0.256	-0.019	0.281	-0.011	0.256
0.003	0.045	0.003	0.046	0.003	0.045	-0.005	0.262	-0.007	0.274	-0.005	0.262
-0.004	0.047	-0.004	0.049	-0.004	0.047	0.002	0.250	-0.007	0.266	0.002	0.252
-0.001	0.045	-0.001	0.046	-0.001	0.045	-0.017	0.239	-0.009	0.247	-0.017	0.239
-0.000	0.048	-0.001	0.050	-0.000	0.048	-0.022	0.260	-0.019	0.270	-0.022	0.260
-0.004	0.043	-0.003	0.044	-0.004	0.043	-0.011	0.241	-0.013	0.261	-0.010	0.241
-0.001	0.048	-0.001	0.049	-0.001	0.048	-0.020	0.256	-0.019	0.259	-0.020	0.256
-0.003	0.047	-0.003	0.049	-0.003	0.047	-0.020	0.252	-0.022	0.269	-0.020	0.252

Table 2: Coordinate-wise bias, ESE, ASE and CovP of EASE, obtained using $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$ or $\mathcal{T} := \text{KM}$, for estimating θ_0 under the linear and NL2C models with $p = 10$ and $\mathbf{b}_p = \mathbf{b}_p^{(1)}$. Shown also are the corresponding bias and ESE of the OLS, as well as the DRESS₁, DRESS₃ and MSSL estimators.

6.2. *Application to EMR Data.* We applied our proposed SS procedures to an EMR study of rheumatoid arthritis (RA), a systemic auto-immune disease (AD), conducted at the Partners HealthCare (Liao et al., 2010). The study cohort consists of 3854 RA patients with blood samples stored. The outcome of interest is the (logarithm of) anti-CCP (antibodies to cyclic citrullinated polypeptide), a biomarker that is often used to determine subtypes of RA. Due to cost constraints, anti-CCP was measured only for a random subset of $n = 355$ patients, thereby leading to a SS set-up. To investigate the validity of the MCAR assumption, we report in Table II in the [Supplementary Material](#) summary measures of the distributions in the labeled and unlabeled data for each of the predictors, as well as p-values from various tests for assessing equality of those distributions. The results suggest that the MCAR assumption is appropriate in this study.

We relate the log anti-CCP level to a set of $p = 24$ clinical variables \mathbf{X} related to ADs, including age, gender, race; total counts of codified and/or narrative mentions extracted from physicians' notes via natural language processing (NLP) for various RA related conditions including RA, Lupus, Polymyalgiarheumatica (PmR), Spondyloarthritis (SpA), as well as various RA medications; indicators of seropositivity and radiological evidence of erosion; mentions of rheumatoid factor (RF), as well as anti-CCP positivity from prior medical history. Since the tests for RF and anti-CCP were not always ordered, missing indicators for these variables were also included. All count variables were transformed as: $x \rightarrow \log(1 + x)$ to increase stability of the model fitting. All predictors were normalized to have unit variance.

We obtained the OLS, the EASE using both $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$ and $\mathcal{T} := \text{KM}$ in the smoothing step, as well as the DRESS_1 estimator for comparison. For EASE, we again used $\mathbb{K} = 5$ and for the $\text{KS}_{2, \mathbf{P}_2}$ smoother, \mathbf{P}_2 was obtained using SIR with $H = 80$ slices of equal width. In Table 3, we present the coordinate-wise estimates of the regression parameters along with their estimated SEs and the corresponding p-values based on these estimates. Overall, the point estimators from all methods are quite close to each other. Our proposed EASE, with both KS and KM smoothers, is substantially more efficient than the OLS across all coordinates with efficiency ranging from about 1.4 to 2.4. The DRESS_1 estimator improved estimation for a few coordinates but the efficiency remains comparable to the OLS for most coordinates. This again suggests the advantage of our proposed estimators compared to both OLS and other SS estimators.

Predictors	OLS ($\hat{\theta}$)			EASE ($\text{KS}_{2, \mathbf{P}_2}$)				RE	DRESS ₁				EASE (KM)			
	Est	SE	Pval	Est	SE	Pval	RE		Est	SE	Pval	RE	Est	SE	Pval	RE
Age	.105	.076	.168	.106	.064	.099	1.40	.094	.073	.199	1.09	.104	.064	.103	1.42	
Gender	-.032	.059	.589	-.028	.050	.570	1.41	-.027	.058	.638	1.04	-.031	.049	.524	1.44	
Race	-.041	.065	.534	-.042	.055	.452	1.40	-.044	.067	.511	.95	-.040	.055	.462	1.41	
Lupus	.038	.066	.563	.048	.052	.359	1.59	.021	.063	.731	1.11	.037	.051	.464	1.70	
PmR	-.075	.044	.088	-.076	.031	.013	2.07	-.074	.031	.016	2.10	-.075	.030	.014	2.04	
RA	.015	.089	.862	.012	.076	.879	1.37	.008	.080	.923	1.23	.016	.075	.832	1.30	
SpA	-.137	.102	.177	-.133	.072	.063	2.02	-.128	.075	.089	1.82	-.136	.066	.038	2.37	
Other ADs	-.022	.078	.775	-.024	.058	.679	1.79	-.018	.067	.792	1.35	-.022	.056	.692	1.93	
Erosion	.076	.070	.278	.078	.059	.184	1.44	.085	.069	.221	1.03	.076	.058	.189	1.47	
Seropositivity	.056	.062	.370	.054	.053	.310	1.37	.041	.061	.496	1.05	.055	.052	.296	1.41	
Anti-CCP _{prior}	.572	.136	.000	.557	.110	.000	1.54	.567	.123	.000	1.23	.568	.107	.000	1.60	
Anti-CCP _{miss}	.527	.123	.000	.520	.097	.000	1.61	.508	.115	.000	1.15	.523	.096	.000	1.64	
RF	.128	.081	.113	.125	.066	.059	1.49	.149	.079	.059	1.05	.127	.066	.054	1.49	
RF _{miss}	.085	.085	.316	.084	.070	.233	1.46	.137	.080	.088	1.12	.084	.070	.231	1.48	
Azathioprine	-.080	.071	.263	-.074	.056	.185	1.62	-.075	.062	.225	1.33	-.079	.053	.132	1.83	
Enbrel	.138	.070	.048	.133	.058	.021	1.48	.136	.073	.064	.91	.137	.057	.017	1.49	
Gold salts	.138	.050	.006	.136	.043	.002	1.37	.147	.050	.003	1.01	.137	.042	.001	1.40	
Humira	-.051	.068	.453	-.049	.057	.391	1.43	-.057	.067	.389	1.03	-.051	.056	.360	1.49	
Infliximab	.003	.069	.968	.008	.057	.887	1.50	.000	.067	.994	1.07	.003	.055	.959	1.57	
Leftunomide	-.027	.069	.697	-.023	.058	.693	1.40	-.031	.071	.660	.93	-.026	.057	.644	1.45	
Methotrexate	-.021	.073	.775	-.024	.061	.699	1.42	-.025	.073	.728	1.01	-.022	.060	.720	1.46	
Plaquenil	-.043	.069	.540	-.038	.057	.503	1.47	-.044	.070	.532	.98	-.042	.057	.460	1.48	
Sulfasalazine	-.114	.074	.125	-.116	.063	.064	1.39	-.105	.072	.145	1.06	-.113	.061	.065	1.45	
Other meds.	-.042	.074	.570	-.052	.060	.385	1.52	-.052	.071	.466	1.10	-.042	.059	.473	1.59	

Table 3: Estimates (Est) of the regression coefficients based on OLS, EASE obtained using either $\mathcal{T} := \text{KS}_{2, \mathbf{P}_2}$ or $\mathcal{T} := \text{KM}$, as well as DRESS₁, along with their estimated standard errors (SE) and the corresponding p-values (Pval.) for testing the null effect of each predictor. Shown also are the relative efficiencies (RE) of all the estimators compared to the OLS.

We also estimated the prediction errors (PEs) for each of the fitted linear models based on the aforementioned estimation methods via CV. To remove potential randomness in the CV partitions, we averaged over 10 replications of leave-5-out CV estimates. The PE was about 1.28 for EASE with both smoothers, 1.29 for OLS and 1.30 for DRESS₁. For prediction purposes, we may also directly employ non-parametric estimates of the conditional mean rather than the fitted linear models. The PE in fact is slightly larger when we use $\hat{m}(\mathbf{x}; \hat{\mathbf{P}}_r)$ or $\hat{\mu}(\mathbf{x}; \hat{\mathcal{P}}_{r, \mathbb{K}})$. The PE was 1.34 for KS and 1.33 for KM based on $\hat{m}(\mathbf{x}; \hat{\mathbf{P}}_r)$, and 1.30 for KS and 1.28 for KM based on $\hat{\mu}(\mathbf{x}; \hat{\mathcal{P}}_{r, \mathbb{K}})$. This confirms that while the linear model may be mis-specified, it may often be preferable to non-parametric models in practice as it may achieve simplicity without substantial loss in prediction performance.

7. Discussion. We have developed in this paper an efficient and adaptive estimation strategy for the SS linear regression problem. The adaptive property possessed by the proposed EASE is crucial for advocating ‘safe’ use of the unlabeled data and is often unaddressed in the existing literature. In general, the magnitude of the efficiency gain with EASE depends on the inherent degree of non-linearity in $\mathbb{E}(Y|\mathbf{X})$ and the extent of sufficiency of the

underlying SNP imputation. In particular, if the imputation is sufficient or the working linear model is correct, $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ is further optimal among a wide class of estimators. We obtained theoretical results along with IF expansions for $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ and $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}^E$ substantiating all our claims and also validated them based on numerical studies. The double CV method further facilitates accurate inference, overcoming potential over-fitting issues in finite samples due to smoothing. An R code for implementing EASE is available upon request.

The proposed SNP imputation, the key component of EASE, apart from being flexible and scalable, enjoys several useful properties. The refitting step and CV play a crucial role in reducing the bias of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$, and for $\mathcal{T} := \text{KS}$ in particular, eradicate any under-smoothing or higher order kernel requirements: two undesirable, yet often inevitable, conditions required for $n^{\frac{1}{2}}$ -consistency of two-step estimators based on a first step of smoothing. Theorem 4.2, apart from showing the distinct behaviour of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, 1)}$ compared to $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ for $\mathbb{K} \geq 2$, also highlights the key role of CV in completely removing kernel order restrictions, apart from addressing over-fitting issues. The error rates in the results of Theorems 4.1-4.2 are quite sharp and account for any estimation error from $\hat{\mathbf{P}}_r$. The regularity conditions required are also fairly mild and standard in the literature. The continuity assumption on \mathbf{X} in Sections 3.1 and 4 is mostly for the convenience of proofs, and the results continue to hold for more general \mathbf{X} . Lastly, while we have focussed here on linear regression for simplicity, our methods can indeed be easily adapted to other regression problems such as logistic regression for binary outcomes.

When the goal is solely that of prediction, one obviously does not have to employ linear regression models, and models that incorporate non-linear effects can be helpful. For such settings, the estimators $\hat{m}(\mathbf{x}; \hat{\mathbf{P}}_r)$ or $\hat{\mu}(\mathbf{x}; \hat{\mathbf{P}}_{r, \mathbb{K}})$, obtained as by-products of our SNP imputation, can themselves serve as potentially useful non-linear predictors. These SNP estimators may substantially outperform naive non-parametric estimators such as a p -dimensional KS estimator, as demonstrated in Table III of the [Supplementary Material](#) for the models considered in our simulation studies. In practice, when the covariates are substantially correlated and the dimension of p is not small as in the EMR example, it is unclear whether non-linear models necessarily provide better prediction performance than the linear models. Under such settings, the linear model also has a clear advantage due to its simplicity. Furthermore, while prediction is a vitally important goal of predictive modeling, association analysis under interpretable models is key to clinical studies for discovery research and efficient estimation of the corresponding model parameters remains an important task.

We end with a comment on the *choice of* $\mathbb{K} \geq 2$ in $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$. While (3.11)

holds for any $\mathbb{K} \geq 2$, the error term in (3.11) depends on \mathbb{K} through $c_{n_{\mathbb{K}}}^-$ and more precisely, through $\tilde{c}_{n_{\mathbb{K}}}^- = \mathbb{K}^{\frac{1}{2}} c_{n_{\mathbb{K}}}^-$. Since \mathbb{K} is fixed, $c_{n_{\mathbb{K}}}^-$ and $\tilde{c}_{n_{\mathbb{K}}}^-$ are asymptotically equivalent. But for a given n , $c_{n_{\mathbb{K}}}^-$ is expected to decrease with \mathbb{K} , while $\tilde{c}_{n_{\mathbb{K}}}^-$ is likely to increase. It is however desirable that both are small since $c_{n_{\mathbb{K}}}^-$ inherently controls the efficiency of the SNP imputation, while $\tilde{c}_{n_{\mathbb{K}}}^-$ directly controls the bias of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$. Hence, a reasonable choice of $\mathbb{K} \geq 2$ may be based on minimizing: $(c_{n_{\mathbb{K}}}^{-2} + \lambda \tilde{c}_{n_{\mathbb{K}}}^{-2})$ for some $\lambda \geq 0$. Since the (first order) asymptotic variance of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ is independent of \mathbb{K} , this is equivalent to a penalized minimization of the asymptotic MSE of $\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})}$ with λ denoting the weightage of the (lower order) bias relative to the (first order) variance. In general, the optimal \mathbb{K} should be inversely related to λ . Conversely, choice of any \mathbb{K} may be viewed to have an associated regularization effect (through λ) resulting in a ‘variance-bias trade-off’ with smaller \mathbb{K} leading to lower bias at the cost of some efficiency, and higher \mathbb{K} leading to improved efficiency in lieu of some bias. In practice, we find that $\mathbb{K} = 5$ works well, and $\mathbb{K} = 10$ tends to give slightly smaller MSE at the cost of increased bias.

APPENDIX A

A.1. Preliminaries. The following Lemmas A.1-A.3 would be useful in the proofs of the main theorems. The proofs of these lemmas as well as Theorems 3.1, 4.1 and 4.2 can be found in the [Supplementary Material](#).

LEMMA A.1. *Let $\mathbf{Z} \in \mathbb{R}^l$ be any random vector and $\mathbf{g}(\mathbf{Z}) \in \mathbb{R}^d$ be any measurable function of \mathbf{Z} , where l and d are fixed. Let $\mathbb{S}_n = \{\mathbf{Z}_i\}_{i=1}^n \perp\!\!\!\perp \mathbb{S}_m = \{\mathbf{Z}_j\}_{j=1}^m$ be two random samples of n and m i.i.d. observations of \mathbf{Z} respectively. Let $\hat{\mathbf{g}}_n(\cdot)$ be any estimator of $\mathbf{g}(\cdot)$ based on \mathbb{S}_n such that the random sequence: $\hat{T}_n \equiv \sup_{\mathbf{z} \in \chi} \|\hat{\mathbf{g}}_n(\mathbf{z})\|$ is $O_p(1)$, where $\chi \subseteq \mathbb{R}^l$ denotes the support of \mathbf{Z} . Let $\hat{\mathbf{G}}_{n,m}$ denote the (double) random sequence: $m^{-1} \sum_{\mathbf{Z}_j \in \mathbb{S}_m} \hat{\mathbf{g}}_n(\mathbf{Z}_j)$, and let $\bar{\mathbf{G}}_n$ denote the random sequence: $\mathbb{E}_{\mathbb{S}_m}(\hat{\mathbf{G}}_{n,m}) = \mathbb{E}_{\mathbf{Z}}\{\hat{\mathbf{g}}_n(\mathbf{Z})\}$, where $\mathbb{E}_{\mathbf{Z}}(\cdot)$ denotes expectation w.r.t. $\mathbf{Z} \in \mathbb{S}_m \perp\!\!\!\perp \mathbb{S}_n$, and all expectations involved are assumed to be finite almost surely (a.s.) $[\mathbb{S}_n] \forall n$. Then: (a) $\hat{\mathbf{G}}_{n,m} - \bar{\mathbf{G}}_n = O_p(m^{-\frac{1}{2}})$, and (b) as long as $\mathbf{g}(\cdot)$ has finite 2nd moments, $m^{-1} \sum_{\mathbf{Z}_j \in \mathbb{S}_m} \mathbf{g}(\mathbf{Z}_j) - \mathbb{E}_{\mathbf{Z}}\{\mathbf{g}(\mathbf{Z})\} = O_p(m^{-\frac{1}{2}})$.*

The next two lemmas would be useful in the proof of Theorem 4.2. They may also be of more general use in other applications that involve controlling empirical process terms indexed by KS estimators. Suppose that Assumption 2.1 (a) holds, and consider the KS framework introduced in Section 4. Let

$l_{\mathbf{P}_r}(\mathbf{w}) = m_{\mathbf{P}_r}(\mathbf{w})f_{\mathbf{P}_r}(\mathbf{w})$ and $\tilde{\varphi}_{\mathbf{P}_r}^{(\varrho)}(\mathbf{w}) = (nh^r)^{-1} \sum_{i=1}^n K_h(\mathbf{w}, \mathbf{P}'_r \mathbf{X}_i) Y_i^\varrho$, for $\varrho = 0, 1$. Let $\tilde{f}_{\mathbf{P}_r} = \tilde{\varphi}_{\mathbf{P}_r}^{(0)}$, $\tilde{l}_{\mathbf{P}_r} = \tilde{\varphi}_{\mathbf{P}_r}^{(1)}$, $\tilde{m}_{\mathbf{P}_r} = \tilde{l}_{\mathbf{P}_r} / \tilde{f}_{\mathbf{P}_r}$, $\varphi_{\mathbf{P}_r}^{(0)} = f_{\mathbf{P}_r}$ and $\varphi_{\mathbf{P}_r}^{(1)} = l_{\mathbf{P}_r}$. Let $\varphi^{(\varrho)}(\mathbf{x}; \mathbf{P}_r) = \varphi_{\mathbf{P}_r}^{(\varrho)}(\mathbf{P}'_r \mathbf{x})$ and $\tilde{\varphi}^{(\varrho)}(\mathbf{x}; \mathbf{P}_r) = \tilde{\varphi}_{\mathbf{P}_r}^{(\varrho)}(\mathbf{P}'_r \mathbf{x}) \forall \varrho = 0, 1$. Let $\tilde{f} = \tilde{\varphi}^{(0)}$, $\tilde{l} = \tilde{\varphi}^{(1)}$ and $\tilde{m} = \tilde{l} / \tilde{f}$. Now, let \mathbb{P}_n denote the empirical probability measure on \mathbb{R}^p based on $\{\mathbf{X}_i\}_{i=1}^n$, and for any measurable function $\gamma(\cdot)$ (possibly vector valued) of \mathbf{X} , let $\mathbb{G}_n^*(\gamma) = n^{\frac{1}{2}} \int \gamma(\mathbf{x})(\mathbb{P}_n - \mathbb{P}_{\mathbf{X}})(d\mathbf{x})$.

LEMMA A.2. *Consider the set-up introduced above. For any fixed integer $d \geq 1$, let $\lambda(\cdot)$ be any \mathbb{R}^d -valued measurable function of \mathbf{X} that is bounded a.s. $[\mathbb{P}_{\mathbf{X}}]$. Define: $b_n^{(1)} = n^{-\frac{1}{2}}h^{-r} + h^q$ and $a_{n,2} = (\log n)^{\frac{1}{2}}(nh^r)^{-\frac{1}{2}} + h^q$. Assume $b_n^{(1)} = o(1)$ for (A.1) and $n^{\frac{1}{2}}a_{n,2}^2 = o(1)$ for (A.2) below. Then, under Assumption 4.1 (i)-(v), and $\forall \varrho \in \{0, 1\}$,*

$$(A.1) \quad \mathbb{G}_n^*[\lambda(\cdot)\{\tilde{\varphi}^{(\varrho)}(\cdot; \mathbf{P}_r) - \varphi^{(\varrho)}(\cdot; \mathbf{P}_r)\}] = O_p(b_n^{(1)}) = o_p(1), \quad \text{and}$$

$$(A.2) \quad \mathbb{G}_n^*[\lambda(\cdot)\{\tilde{m}(\cdot; \mathbf{P}_r) - m(\cdot; \mathbf{P}_r)\}] = O_p(n^{\frac{1}{2}}a_{n,2}^2) = o_p(1).$$

Let $\hat{\varphi}^{(\varrho)}(\mathbf{x}; \hat{\mathbf{P}}_r) = (nh^r)^{-1} \sum_{i=1}^n K_h(\hat{\mathbf{P}}'_r \mathbf{x}, \hat{\mathbf{P}}'_r \mathbf{X}_i) Y_i^\varrho \forall \varrho \in \{0, 1\}$, where $\hat{\mathbf{P}}_r$ is as in section 3.2 and all other notations are the same as in the set-up of Lemma A.2. Let $\hat{f}(\mathbf{x}; \hat{\mathbf{P}}_r) = \hat{\varphi}^{(0)}(\mathbf{x}; \hat{\mathbf{P}}_r)$ and $\hat{l}(\mathbf{x}; \hat{\mathbf{P}}_r) = \hat{\varphi}^{(1)}(\mathbf{x}; \hat{\mathbf{P}}_r)$. Then:

LEMMA A.3. *Consider the set-up of Lemma A.2. Let $\hat{\varphi}^{(\varrho)}(\mathbf{x}; \hat{\mathbf{P}}_r)$ be as above, and let $\lambda(\cdot)$ be as in Lemma A.2. Suppose $(\hat{\mathbf{P}}_r - \mathbf{P}_r) = O_p(\alpha_n)$ for some $\alpha_n = o(1)$. Assume $b_n^{(2)} = o(1)$, where $b_n^{(2)} = \alpha_n + n^{-\frac{1}{2}}\alpha_n h^{-(r+1)} + n^{\frac{1}{2}}\alpha_n^2(h^{-2} + n^{-1}h^{-(r+2)})$. Then, under Assumption 4.1,*

$$(A.3) \quad \mathbb{G}_n^*[\lambda(\cdot)\{\hat{\varphi}^{(\varrho)}(\cdot; \hat{\mathbf{P}}_r) - \tilde{\varphi}^{(\varrho)}(\cdot; \mathbf{P}_r)\}] = O_p(b_n^{(2)}) = o_p(1) \quad \forall \varrho \in \{0, 1\}.$$

A.2. Proof of Theorem 3.2. Let $\mathbf{\Gamma}_n = \frac{1}{n} \sum_{i=1}^n \vec{\mathbf{X}}_i \vec{\mathbf{X}}_i'$, and

$$\mathbf{T}_n^{(1)} = \frac{1}{n} \sum_{i=1}^n \vec{\mathbf{X}}_i \{Y_i - \mu(\mathbf{X}_i; \mathbf{P}_r)\}, \quad \mathbf{T}_{n, \mathbb{K}}^{(2)} = \frac{1}{n} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \vec{\mathbf{X}}_i \hat{\Delta}_k(\mathbf{X}_i; \mathbf{P}_r, \hat{\mathbf{P}}_{r,k}).$$

Then, using (3.4)-(3.8), it is straightforward to see that:

$$(A.4) \quad \mathbb{E}[\vec{\mathbf{X}}\{Y - \mu(\mathbf{X}; \mathbf{P}_r)\}] \equiv \mathbb{E}[\vec{\mathbf{X}}\{Y - m(\mathbf{X}; \mathbf{P}_r) - \vec{\mathbf{X}}'\boldsymbol{\eta}_{\mathbf{P}_r}\}] = \mathbf{0}, \quad \text{and}$$

$$(A.5) \quad \mathbf{\Gamma}_n \left(\hat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r} \right) = \mathbf{T}_n^{(1)} - \mathbf{T}_{n, \mathbb{K}}^{(2)}.$$

Under (A.4), Assumptions 2.1 (a) and (i), it follows from Lemma A.1 (b) that $\mathbf{T}_n^{(1)} = O_p(n^{-\frac{1}{2}})$. Next, due to assumption (ii) and boundedness of \mathbf{X} ,

$$\|\mathbf{T}_{n,\mathbb{K}}^{(2)}\| \leq n^{-1} \sum_{k=1}^{\mathbb{K}} \sum_{i \in \mathcal{I}_k} \sup_{\mathbf{x} \in \mathcal{X}} \{\|\vec{\mathbf{x}}\| |\hat{\Delta}_k(\mathbf{x}; \mathbf{P}_r, \hat{\mathbf{P}}_{r,k})|\} = O_p(c_{n_{\mathbb{K}}}^-).$$

Finally, under Assumption 2.1 (a), we have: $\mathbf{\Gamma}_n = \mathbf{\Gamma} + O_p(n^{-\frac{1}{2}})$ using Lemma A.1 (b). Further, since $\mathbf{\Gamma}_n$ is invertible a.s., $\mathbf{\Gamma}_n^{-1} = \mathbf{\Gamma}^{-1} + O_p(n^{-\frac{1}{2}})$. Using all these facts, we then have: $(\hat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r}) = \mathbf{\Gamma}_n^{-1}(\mathbf{T}_n^{(1)} - \mathbf{T}_{n,\mathbb{K}}^{(2)}) = \mathbf{\Gamma}^{-1}(\mathbf{T}_n^{(1)} - \mathbf{T}_{n,\mathbb{K}}^{(2)}) + O_p\{n^{-\frac{1}{2}}(n^{-\frac{1}{2}} + c_{n_{\mathbb{K}}}^-)\}$. Thus,

$$(A.6) \quad (\hat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r}) = \mathbf{\Gamma}^{-1}(\mathbf{T}_n^{(1)} - \mathbf{T}_{n,\mathbb{K}}^{(2)}) + O_p(n^{-1} + n^{-\frac{1}{2}}c_{n_{\mathbb{K}}}^-).$$

Next, let us define:

$$\begin{aligned} \mathbf{\Gamma}_N &= N^{-1} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \vec{\mathbf{X}}_j', \quad \mathbf{R}_N^{(1)} = N^{-1} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \{\mu(\mathbf{X}_j; \mathbf{P}_r) - \vec{\mathbf{X}}_j' \boldsymbol{\theta}_0\}, \\ \text{and } \hat{\mathbf{R}}_{N,n}^{(\mathbb{K})} &= N^{-1} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \{\hat{\mu}(\mathbf{X}_j; \hat{\mathcal{P}}_{r,\mathbb{K}}) - \mu(\mathbf{X}_j; \mathbf{P}_r)\}. \end{aligned}$$

Then, using (3.7), we have:

$$\mathbf{\Gamma}_N(\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0) = N^{-1} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j [\hat{\mu}(\mathbf{X}_j; \hat{\mathcal{P}}_{r,\mathbb{K}}) - \vec{\mathbf{X}}_j' \boldsymbol{\theta}_0] = \mathbf{R}_N^{(1)} + \hat{\mathbf{R}}_{N,n}^{(\mathbb{K})}.$$

Next, using (3.4)-(3.8), we have: $\hat{\mathbf{R}}_{N,n}^{(\mathbb{K})} = \mathbf{\Gamma}_N(\hat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r}) + \hat{\mathbf{S}}_{N,n}^{(\mathbb{K})}$, where

$$\hat{\mathbf{S}}_{N,n}^{(\mathbb{K})} = \mathbb{K}^{-1} \sum_{k=1}^{\mathbb{K}} \{N^{-1} \sum_{j=n+1}^{n+N} \vec{\mathbf{X}}_j \hat{\Delta}_k(\mathbf{X}_j; \mathbf{P}_r, \hat{\mathbf{P}}_{r,k})\}.$$

Hence, we have: $\mathbf{\Gamma}_N(\hat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0) = \mathbf{\Gamma}_N(\hat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r}) + \mathbf{R}_N^{(1)} + \hat{\mathbf{S}}_{N,n}^{(\mathbb{K})}$.

Now, under assumptions (i)-(ii) and Assumption 2.1 (a), we have:

$$(I) \quad \sum_{k=1}^{\mathbb{K}} \sup_{\mathbf{x} \in \mathcal{X}} \|\vec{\mathbf{x}} \hat{\Delta}_k(\mathbf{x}; \mathbf{P}_r, \hat{\mathbf{P}}_{r,k})\| = O_p(1),$$

so that using Lemma A.1 (a), $\hat{\mathbf{S}}_{N,n}^{(\mathbb{K})} = \mathbb{K}^{-1} \sum_{k=1}^{\mathbb{K}} \hat{\mathbf{S}}_{n,k}^* + O_p(N^{-\frac{1}{2}})$, where $\hat{\mathbf{S}}_{n,k}^* = \mathbb{E}_{\mathbf{X}}\{\vec{\mathbf{X}} \hat{\Delta}_k(\mathbf{X}; \mathbf{P}_r, \hat{\mathbf{P}}_{r,k})\} \quad \forall 1 \leq k \leq \mathbb{K}$;

$$(II) \quad \mathbf{R}_N^{(1)} = \mathbb{E}[\vec{\mathbf{X}} \{\mu(\mathbf{X}; \mathbf{P}_r) - \vec{\mathbf{X}}' \boldsymbol{\theta}_0\}] + O_p(N^{-\frac{1}{2}}) = O_p(N^{-\frac{1}{2}})$$

from Lemma A.1 (b) and $\mathbb{E}[\vec{\mathbf{X}}\{\mu(\mathbf{X}; \mathbf{P}_r) - \vec{\mathbf{X}}'\boldsymbol{\theta}_0\}] = \mathbf{0}$ due to (A.4) and 2.1; and lastly, (III) $\boldsymbol{\Gamma}_N^{-1} = \boldsymbol{\Gamma}^{-1} + O_p(N^{-\frac{1}{2}})$. It then follows from (I)-(III) that

$$(A.7) \quad \widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0 = (\widehat{\boldsymbol{\eta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\eta}_{\mathbf{P}_r}) + \mathbb{K}^{-1}\boldsymbol{\Gamma}^{-1} \sum_{k=1}^{\mathbb{K}} \widehat{\mathbf{S}}_{n,k}^* + O_p(N^{-\frac{1}{2}}).$$

Using (A.6) and (3.9) in (A.7), we then have:

$$\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0 = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\psi}(\mathbf{Z}_i; \mathbf{P}_r) - \boldsymbol{\Gamma}^{-1} \frac{1}{\mathbb{K}} \sum_{k=1}^{\mathbb{K}} \left\{ \frac{1}{n_{\mathbb{K}}} \sum_{i \in \mathcal{I}_k} \widehat{\mathbf{G}}_k(\mathbf{X}_i) \right\} + O_p(b_{n, \mathbb{K}}),$$

where $b_{n, \mathbb{K}} = n^{-1} + n^{-\frac{1}{2}}c_{n_{\mathbb{K}}}^- + N^{-\frac{1}{2}}$. It follows, as claimed in (3.10), that

$$(A.8) \quad n^{\frac{1}{2}}(\widehat{\boldsymbol{\theta}}_{(\mathbf{P}_r, \mathbb{K})} - \boldsymbol{\theta}_0) = n^{-\frac{1}{2}} \sum_{i=1}^n \boldsymbol{\psi}(\mathbf{Z}_i; \mathbf{P}_r) - \boldsymbol{\Gamma}^{-1} \mathbb{G}_{n, \mathbb{K}} + O_p(c_{n, \mathbb{K}}^*) \quad \blacksquare$$

We next show that $\mathbb{G}_{n, \mathbb{K}} = O_p(c_{n_{\mathbb{K}}}^-)$ for any fixed $\mathbb{K} \geq 2$. To this end, let $\mathbb{T}_k^{(n)} = (n_{\mathbb{K}})^{-\frac{1}{2}} \sum_{i \in \mathcal{I}_k} \widehat{\mathbf{G}}_k(\mathbf{X}_i)$, $\widehat{D}_k = \sup_{\mathbf{x} \in \mathcal{X}} |\widehat{\Delta}_k(\mathbf{x}; \mathbf{P}_r, \widehat{\mathbf{P}}_{r, k})|$ and $C = \sup_{\mathbf{x} \in \mathcal{X}} \|\vec{\mathbf{X}}\| < \infty$. For any subset $\mathcal{A} \subseteq \mathcal{L}$, let $\mathbb{P}_{\mathcal{A}}$ denote the joint distribution of the observations in \mathcal{A} , and let $\mathbb{E}_{\mathcal{A}}(\cdot)$ denote expectation w.r.t. $\mathbb{P}_{\mathcal{A}}$. By definition, $\mathbb{G}_{n, \mathbb{K}} = \mathbb{K}^{-\frac{1}{2}} \sum_{k=1}^{\mathbb{K}} \mathbb{T}_k^{(n)} = O_p(c_{n_{\mathbb{K}}}^-)$ if and only if given any $\epsilon > 0$, $\exists M_\epsilon > 0$ such that $\mathbb{P} \left(\|\mathbb{G}_{n, \mathbb{K}}\| > M_\epsilon c_{n_{\mathbb{K}}}^- \right) \leq \epsilon \forall n$. Note that for any $M > 0$,

$$(A.9) \quad \begin{aligned} \mathbb{P} \left(\|\mathbb{G}_{n, \mathbb{K}}\| > M c_{n_{\mathbb{K}}}^- \right) &\leq \mathbb{P} \left(\mathbb{K}^{-\frac{1}{2}} \sum_{k=1}^{\mathbb{K}} \|\mathbb{T}_k^{(n)}\| > M c_{n_{\mathbb{K}}}^- \right) \\ &\leq \sum_{k=1}^{\mathbb{K}} \mathbb{P} \left(\mathbb{K}^{-\frac{1}{2}} \|\mathbb{T}_k^{(n)}\| > \frac{M c_{n_{\mathbb{K}}}^-}{\mathbb{K}} \right) \leq \sum_{k=1}^{\mathbb{K}} \sum_{l=1}^{p+1} \mathbb{P} \left\{ |\mathbb{T}_{k[l]}^{(n)}| > \frac{M c_{n_{\mathbb{K}}}^-}{\mathbb{K}^{\frac{1}{2}} (p+1)^{\frac{1}{2}}} \right\} \\ &= \sum_{k=1}^{\mathbb{K}} \sum_{l=1}^{p+1} \mathbb{E}_{\mathcal{L}_k^-} \left[\mathbb{P}_{\mathcal{L}_k} \left\{ |\mathbb{T}_{k[l]}^{(n)}| > \frac{M c_{n_{\mathbb{K}}}^-}{\mathbb{K}^{\frac{1}{2}} (p+1)^{\frac{1}{2}}} \mid \mathcal{L}_k^- \right\} \right], \end{aligned}$$

where the steps follow from repeated use of Bonferroni's inequality and other standard arguments. Now, conditional on \mathcal{L}_k^- ($\perp \mathcal{L}_k$, with $\mathbb{K} \geq 2$), $n_{\mathbb{K}}^{-\frac{1}{2}} \mathbb{T}_k^{(n)}$ is a centered sum of the i.i.d. random vectors $\{\vec{\mathbf{X}}_i \widehat{\Delta}_k(\mathbf{X}_i; \mathbf{P}_r, \widehat{\mathbf{P}}_{r, k})\}_{i \in \mathcal{I}_k}$ which, due to assumption (ii) and the compactness of \mathcal{X} , are bounded by: $C \widehat{D}_k < \infty$

a.s. $[\mathbb{P}_{\mathcal{L}_k^-}] \forall k, n$. Hence, applying Hoeffding's inequality to $\mathbb{T}_{k[l]}^{(n)} \forall l$, we have:

$$(A.10) \quad \mathbb{P}_{\mathcal{L}_k} \left\{ \left| \mathbb{T}_{k[l]}^{(n)} \right| > \frac{M c_{n_{\mathbb{K}}}^-}{\mathbb{K}^{\frac{1}{2}} (p+1)^{\frac{1}{2}}} \mid \mathcal{L}_k^- \right\} \leq 2 \exp \left\{ - \frac{M^2 c_{n_{\mathbb{K}}}^2}{2(p+1) \mathbb{K} C^2 \widehat{D}_k^2} \right\}$$

a.s. $[\mathbb{P}_{\mathcal{L}_k^-}] \forall n$; for each $k \in \{1, \dots, \mathbb{K}\}$ and $\forall 1 \leq l \leq (p+1)$.

Now, since $\widehat{D}_k = O_p(c_{n_{\mathbb{K}}}^-)$, $(c_{n_{\mathbb{K}}}^- / \widehat{D}_k) \geq 0$ is stochastically bounded away from 0. Thus, $\forall k$, and for any given $\epsilon > 0$, $\exists \delta(k, \epsilon) > 0$ (independent of n) such that: $\mathbb{P}_{\mathcal{L}_k^-} \{(c_{n_{\mathbb{K}}}^- / \widehat{D}_k) \leq \delta(k, \epsilon)\} \leq \epsilon^* \forall n$, where $\epsilon^* = \epsilon / \{4\mathbb{K}(p+1)\} > 0$. Let $\widetilde{\delta}(\mathbb{K}, \epsilon) = \min\{\delta(k, \epsilon) : k = 1, \dots, \mathbb{K}\} > 0$ (as \mathbb{K} is fixed). Let $\mathbb{A}(k, \epsilon)$ denote the event: $\{(c_{n_{\mathbb{K}}}^- / \widehat{D}_k) \leq \widetilde{\delta}(\mathbb{K}, \epsilon)\}$, and let $\mathbb{A}^c(k, \epsilon)$ be its complement. Then, $\mathbb{P}_{\mathcal{L}_k^-} \{\mathbb{A}(k, \epsilon)\} \leq \epsilon^*$, while on $\mathbb{A}^c(k, \epsilon)$, $(c_{n_{\mathbb{K}}}^- / \widehat{D}_k) > \widetilde{\delta}(\mathbb{K}, \epsilon)$. Thus, the bound in (A.10) is dominated by: $2 \exp[-M^2 \widetilde{\delta}^2(\mathbb{K}, \epsilon) / \{2(p+1) \mathbb{K} C^2\}]$ on $\mathbb{A}^c(k, \epsilon)$, and trivially by 2 on $\mathbb{A}(k, \epsilon) \forall k$. Plugging the bound of (A.10) into (A.9) and using all these facts, we then have:

$$\begin{aligned} \mathbb{P} \left(\|\mathbb{G}_{n, \mathbb{K}}\| > M c_{n_{\mathbb{K}}}^- \right) &\leq \sum_{k=1}^{\mathbb{K}} \sum_{l=1}^{p+1} \mathbb{E}_{\mathcal{L}_k^-} \left[2 \exp \left\{ - \frac{M^2 c_{n_{\mathbb{K}}}^2}{2(p+1) \mathbb{K} C^2 \widehat{D}_k^2} \right\} \right] \\ &= \sum_{k=1}^{\mathbb{K}} \sum_{l=1}^{p+1} \mathbb{E}_{\mathcal{L}_k^-} \left[2 \exp \left\{ - \frac{M^2 c_{n_{\mathbb{K}}}^2}{2(p+1) \mathbb{K} C^2 \widehat{D}_k^2} \right\} \{1_{\mathbb{A}^c(k, \epsilon)} + 1_{\mathbb{A}(k, \epsilon)}\} \right] \\ &\leq \sum_{k=1}^{\mathbb{K}} \sum_{l=1}^{p+1} \left[2 \exp \left\{ - \frac{M^2 \widetilde{\delta}^2(\mathbb{K}, \epsilon)}{2(p+1) \mathbb{K} C^2} \right\} \mathbb{P}_{\mathcal{L}_k^-} \{\mathbb{A}^c(k, \epsilon)\} + 2 \mathbb{P}_{\mathcal{L}_k^-} \{\mathbb{A}(k, \epsilon)\} \right] \\ &\leq 2\mathbb{K}(p+1) \left[\exp \left\{ - \frac{M^2 \widetilde{\delta}^2(\mathbb{K}, \epsilon)}{2(p+1) \mathbb{K} C^2} \right\} + \epsilon^* \right] \\ (A.11) \quad &\leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \quad (\text{with some suitable choice } M_\epsilon \text{ for } M), \end{aligned}$$

where the last step follows from noting the definition of ϵ^* and choosing M_ϵ to be any M large enough such that $4 \exp[-M^2 \widetilde{\delta}^2(\mathbb{K}, \epsilon) / \{2(p+1) \mathbb{K} C^2\}] \leq \epsilon / \{\mathbb{K}(p+1)\}$. Thus, (A.11) shows $\mathbb{G}_{n, \mathbb{K}} = O_p(c_{n_{\mathbb{K}}}^-)$ for any fixed $\mathbb{K} \geq 2$. This further establishes (3.11) and all its associated implications. The proof of Theorem 3.2 is now complete. \blacksquare

ACKNOWLEDGEMENTS

The authors would like to thank Dr. James Robins and Dr. Eric Tchetgen Tchetgen for many helpful discussions throughout the progress of this paper,

as well as the editor Dr. Edward George, the anonymous associate editor and the two referees for their useful comments and suggestions that helped significantly in improving and revising the original version of this article.

SUPPLEMENTARY MATERIAL

Supplement to “Efficient and Adaptive Linear Regression in Semi-Supervised Settings”. The supplement includes: (i) Supplementary results for the simulation studies and the real data analysis, (ii) Brief discussions on generalization of the proposed SS estimators to MAR settings, (iii) Proof of Lemma A.1, (iv) Proof of Theorem 3.1, (v) Proof of Theorem 4.1, and (vi) Proofs of Lemmas A.2-A.3 and Theorem 4.2.

REFERENCES

- ANDREWS, D. W. K. (1995). Nonparametric Kernel Estimation for Semiparametric Models. *Econometric Theory* **11** 560-586.
- BELKIN, M., NIYOGI, P. and SINDHWANI, V. (2006). Manifold Regularization : A Geometric Framework for Learning from Labeled and Unlabeled Examples. *The Journal of Machine Learning Research* **7** 2399-2434.
- CASTELLI, V. and COVER, T. M. (1995). The Exponential Value of Labeled Samples. *Pattern Recognition Letters* **16** 105-111.
- CASTELLI, V. and COVER, T. M. (1996). The Relative Value of Labeled and Unlabeled Samples in Pattern Recognition with an Unknown Mixing Parameter. *IEEE Transactions on Information Theory* **42** 2102-2117.
- CHAPELLE, O., SCHÖLKOPF, B. and ZIEN, A. (2006). *Semi-Supervised Learning*. MIT Press, Cambridge, MA, USA.
- COOK, R. D. (1998). Principal Hessian Directions Revisited (with Discussion). *Journal of the American Statistical Association* **93** 84-100.
- COOK, R. D. and LEE, H. (1999). Dimension Reduction in Binary Response Regression. *Journal of the American Statistical Association* **94** 1187-1200.
- COOK, R. D. and WEISBERG, S. (1991). Discussion of “Sliced Inverse Regression” by K.-C. Li. *Journal of the American Statistical Association* **86** 328-332.
- COZMAN, F. G. and COHEN, I. (2001). Unlabeled Data Can Degrade Classification Performance of Generative Classifiers. Technical Report No. HPL-2001-234, HP Laboratories, Palo Alto, CA, USA.
- COZMAN, F. G., COHEN, I. and CIRELO, M. C. (2003). Semi-Supervised Learning of Mixture Models. In *Proceedings of the Twentieth ICML* 99-106.
- DUAN, N. and LI, K.-C. (1991). Sliced Regression: A Link-Free Regression Method. *The Annals of Statistics* **19** 505-530.
- HANSEN, B. E. (2008). Uniform Convergence Rates for Kernel Estimation with Dependent Data. *Econometric Theory* **24** 726-748.
- KAWAKITA, M. and KANAMORI, T. (2013). Semi-Supervised Learning with Density-Ratio Estimation. *Machine Learning* **91** 189-209.
- KOHANE, I. S. (2011). Using Electronic Health Records to Drive Discovery in Disease Genomics. *Nature Reviews Genetics* **12** 417-428.
- LAFFERTY, J. D. and WASSERMAN, L. (2007). Statistical Analysis of Semi-Supervised Regression. *Advances in Neural Information Processing Systems* **20** 801-808.

- LI, K.-C. (1991). Sliced Inverse Regression for Dimension Reduction. *Journal of the American Statistical Association* **86** 316-327.
- LI, K.-C. (1992). On Principal Hessian Directions for Data Visualization and Dimension Reduction: Another Application of Stein's Lemma. *Journal of the American Statistical Association* **87** 1025-1039.
- LIAO, K. P., CAI, T., GAINER, V. et al. (2010). Electronic Medical Records for Discovery Research in Rheumatoid Arthritis. *Arthritis Care and Research* **62** 1120-1127.
- MASRY, E. (1996). Multivariate Local Polynomial Regression for Time Series: Uniform Strong Consistency and Rates. *Journal of Time Series Analysis* **17** 571-600.
- NEWBY, W. K. (1994). Kernel Estimator of Partial Means and a Generalized Variance Estimator. *Econometric Theory* **10** 1-21.
- NEWBY, W. K., HSIEH, F. and ROBINS, J. (1998). Undersmoothing and Bias Corrected Functional Estimation. Technical Report No. 98-17, Dept. of Economics, MIT, USA.
- NIGAM, K. P. (2001). Using Unlabeled Data to Improve Text Classification. PhD thesis, Carnegie Mellon University, USA. CMU-CS-01-126.
- NIGAM, K., MCCALLUM, A. K., THRUN, S. and MITCHELL, T. (2000). Text Classification from Labeled and Unlabeled Documents Using EM. *Machine Learning* **39** 103-134.
- SEEGER, M. (2002). Learning with Labeled and Unlabeled Data. Technical Report No. EPFL-REPORT-161327, University of Edinburgh, UK.
- SOKOLOVSKA, N., CAPPÉ, O. and YVON, F. (2008). The Asymptotics of Semi-Supervised Learning in Discriminative Probabilistic Models. In *Proceedings of the Twenty Fifth ICML* 984-991.
- ZHANG, T. and OLES, F. J. (2000). The Value of Unlabeled Data for Classification Problems. In *Proceedings of the Seventeenth ICML* 1191-1198.
- ZHU, X. (2005). Semi-Supervised Learning through Graphs. PhD thesis, Carnegie Mellon University, USA. CMU-LTI-05-192.
- ZHU, X. (2008). Semi-Supervised Learning Literature Survey. Technical Report No. 1530, Computer Sciences, University of Wisconsin-Madison, USA.
- ZHU, L.-X. and NG, K. W. (1995). Asymptotics of Sliced Inverse Regression. *Statistica Sinica* **5** 727-736.

ABHISHEK CHAKRABORTTY
DEPARTMENT OF BIostatISTICS
HARVARD UNIVERSITY
BOSTON, MA 02115, USA.
E-MAIL: achakrabortty@mail.harvard.edu

TIANXI CAI
DEPARTMENT OF BIostatISTICS
HARVARD UNIVERSITY
BOSTON, MA 02115, USA.
E-MAIL: tcai@hsph.harvard.edu