

More Efficient Estimation for Logistic Regression with Optimal Subsample

HaiYing Wang

Department of Statistics, University of Connecticut

haiying.wang@uconn.edu

Abstract

Facing large amounts of data, subsampling is a practical technique to extract useful information. For this purpose, Wang et al. (2017) developed an Optimal Subsampling Method under the A-optimality Criterion (OSMAC) for logistic regression that samples more informative data points with higher probabilities. However, the original OSMAC estimator use inverse of optimal subsampling probabilities as weights in the likelihood function. This reduces contributions of more informative data points and the resultant estimator may lose efficiency. In this paper, we propose a more efficient estimator based on OSMAC subsample without weighting the likelihood function. Both asymptotic results and numerical results show that the new estimator is more efficient. In addition, our focus in this paper is inference for the true parameter, while Wang et al. (2017) focuses on approximating the full data estimator. We also develop a new algorithm based on Poisson sampling, which does not require to approximate the optimal subsampling probabilities all at once. This is computationally advantageous when available random-access memory is not enough to hold the full data. Interestingly, asymptotic distributions also show that Poisson sampling produces more efficient estimator if the sampling rate, the ratio of the subsample size to the full data sample size, does not converge to 0. We also obtain the unconditional asymptotic distribution for the estimator based on Poisson sampling.

Keywords: Asymptotic Distribution; Logistic Regression; Massive Data; Optimal Subsampling, Poisson Sampling.

1 Introduction

Extraordinary amounts of data that are collected offer unparalleled opportunities for advancing complicated scientific problems. However, the incredible sizes of big data bring new challenges for data analysis. A major challenge of big data analysis lies with the thirst for computing resources. Faced with this, subsampling has been widely used to reduce the computational burden, in which intended calculations are carried out on a subsample that is drawn from the full data, see Drineas, Kannan and Mahoney (2006*a,b,c*); Mahoney and Drineas (2009); Drineas et al. (2011); Mahoney (2011); Halko et al. (2011); Clarkson and Woodruff (2013); Kleiner et al. (2014); McWilliams et al. (2014); Yang et al. (2016), among others.

A key to success of a subsampling method is to specify nonuniform sampling probabilities so that more informative data points are sampled with higher probabilities. For this purpose, normalized statistical leverage scores or its variants are often used as subsampling probabilities in the context of linear regression, and this approach is termed *algorithmic leveraging* (Ma et al., 2014; Ma and Sun, 2015; Ma et al., 2015). It has demonstrated remarkable performance in making better use of a fixed amount of computing power (Avron et al., 2010; Meng et al., 2014). Statistical leverage scores only contain information in the covariate and do not take into account the information contained in the observed responses. Wang et al. (2017) derived optimal subsampling probabilities that minimize the asymptotic mean squared error (MSE) of the subsampling-based estimator in the context of logistic regression. The optimal subsampling probabilities directly depend on both the covariates and the responses to take more informative subsample. However, Wang et al. (2017) used a weighted maximum likelihood estimator based on the subsample, and more informative data points are assigned smaller weights in the likelihood function. Thus, the resultant estimator may not be efficient.

In this paper, we propose more efficient estimators based on subsample taken randomly according to the optimal subsampling probabilities. Asymptotic distributions will be derived which show that asymptotic variance-covariance matrices of the new estimators are smaller, in Loewner-ordering, than that of the weighted estimator in Wang et al. (2017). We also consider to use Poisson sampling instead of sampling with replacement. Asymptotic distributions show that Poisson sampling is more efficient in parameter estimation when the subsample size is proportional to the full data sample size. It is also computationally beneficial to use Poisson sampling because there is no need to generate a large amount of random numbers all at once.

Before presenting the framework of the paper, we give a brief review of the emerging field

of subsampling-based methods. For linear regression, Drineas, Mahoney and Muthukrishnan (2006) developed a subsampling method and focused on finding influential data units for the least squares (LS) estimates. Drineas et al. (2011) developed an algorithm by processing the data with randomized Hadamard transform and then using uniform subsampling to approximate LS estimates. Drineas et al. (2012) developed an algorithm to approximate statistical leverage scores that are used for algorithmic leveraging. Yang et al. (2015) showed that using normalized square roots of statistical leverage scores as subsampling probabilities yields better approximation than using original statistical leverage scores, if they are very nonuniform. The aforementioned studies focused on developing algorithms for fast approximation of LS estimates. Ma et al. (2014, 2015) and Ma and Sun (2015) considered the statistical properties of algorithmic leveraging. They derived biases and variances of leverage-based subsampling estimators in linear regression and proposed a shrinkage algorithmic leveraging method to improve the performance. Raskutti and Mahoney (2016) considered both the algorithmic and statistical aspects of solving large-scale LS problems using random sketching. Wang et al. (2018) developed an information-based optimal subdata selection method to select subsample deterministically for ordinary LS in linear regression. The aforesaid results were obtained exclusively within the context of linear models. Fithian and Hastie (2014) proposed a computationally efficient local case-control subsampling method for logistic regression with large imbalanced data. Recently, Wang et al. (2017) developed an Optimal Subsampling Method under the A-optimality Criterion (OSMAC) for logistic regression. Although they derived optimal subsampling probabilities, the inference procedure based on weighted likelihood is not efficient.

This paper focuses on logistic regression models, which are widely used for statistical inference in many disciplines, such as business, computer science, education, and genetics, among others (Hosmer Jr et al., 2013). Based on optimal subsample taken according to OSMAC developed in Wang et al. (2017), more efficient methods, in terms of both parameter estimation and numerical computation, will be proposed. The remainder of the paper is organized as follows. Model setups and notations are introduced in Section 2. The OSMAC will also be briefly reviewed in this section. Section 3 presents the more efficient estimator and its asymptotic properties. Section 4 considers Poisson sampling. Section 5 discusses issues related to practical implementation and summaries the methods from Sections 3 and 4 into two practical algorithms. Section 6 gives unconditional asymptotic distributions for the estimator from Poisson sampling. Section 7 evaluates the practical performance of the proposed methods using numerical experiments. Section 8 concludes, and the appendix contains proofs and technical details.

2 Model setup and optimal subsampling

Let $y \in \{0, 1\}$ be a binary response variable and \mathbf{x} be a d dimensional covariate. A logistic regression model describes the conditional probability of $y = 1$ given \mathbf{x} , and it has the following form.

$$\mathbb{P}(y = 1|\mathbf{x}) = p(\boldsymbol{\beta}) = \frac{e^{\mathbf{x}^T \boldsymbol{\beta}}}{1 + e^{\mathbf{x}^T \boldsymbol{\beta}}}, \quad (1)$$

where $\boldsymbol{\beta}$ is a $d \times 1$ vector of unknown regression coefficients belonging to a compact subset of \mathbb{R}^d .

With independent full data of size N from Model (1), say, $\mathcal{D}_N = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$, the unknown parameter $\boldsymbol{\beta}$ is often estimated by the maximum likelihood estimator (MLE), denoted as $\hat{\boldsymbol{\beta}}_f$. It is the maximizer of the log-likelihood function, namely,

$$\hat{\boldsymbol{\beta}}_f = \arg \max_{\boldsymbol{\beta}} \ell_f(\boldsymbol{\beta}) = \arg \max_{\boldsymbol{\beta}} \sum_{i=1}^N \{y_i \mathbf{x}_i^T \boldsymbol{\beta} - \log(1 + e^{\boldsymbol{\beta}^T \mathbf{x}_i})\}. \quad (2)$$

Since there is no general closed-form solution to the MLE, Newton's method or iteratively reweighted least squares method (McCullagh and Nelder, 1989) is often adopted to find it numerically. This typically takes $O(\zeta N d^2)$ time, where ζ is the number of iterations in the optimization procedure (Wang et al., 2017). For super-large data, the computing time $O(\zeta N d^2)$ may be too long to afford, and iterative computation is infeasible if the data volume is larger than the available random-access memory (RAM). To overcome this computational bottleneck for the application of logistic regression on massive data, Wang et al. (2017) developed the OSMAC under the subsampling framework.

Let π_1, \dots, π_N be subsampling probabilities such that $\sum_{i=1}^N \pi_i = 1$. Using subsampling with replacement, draw a random subsample of size n ($\ll N$), according to the probabilities $\{\pi_i\}_{i=1}^N$, from the full data. We use $*$ to indicate any quantity for subsample, namely, denote the covariates, responses, and subsampling probabilities in the subsample as \mathbf{x}_i^* , y_i^* , and π_i^* , respectively, for $i = 1, \dots, n$. Wang et al. (2017) define the subsample estimator $\hat{\boldsymbol{\beta}}_w^\pi$ to be the weighted MLE, i.e.,

$$\hat{\boldsymbol{\beta}}_w^\pi = \arg \max_{\boldsymbol{\beta}} \ell_w^*(\boldsymbol{\beta}) = \arg \max_{\boldsymbol{\beta}} \sum_{i=1}^n \frac{y_i^* \boldsymbol{\beta}^T \mathbf{x}_i^* - \log(1 + e^{\boldsymbol{\beta}^T \mathbf{x}_i^*})}{\pi_i^*}. \quad (3)$$

The key to success here is how to specify the values for π_i 's so that more informative data points are sampled with higher probabilities. Wang et al. (2017) derived optimal subsampling probabilities that minimize the asymptotic MSE of $\hat{\boldsymbol{\beta}}_w^\pi$. They first show that $\hat{\boldsymbol{\beta}}_w^\pi$ is asymptotically normal. Specifically, for large n and N , the conditional distribution of

$\sqrt{n}(\hat{\boldsymbol{\beta}}_w^\pi - \hat{\boldsymbol{\beta}}_f)$ given the full data \mathcal{D}_N can be approximated by a normal distribution with mean $\mathbf{0}$ and variance-covariance matrix $\mathbf{V}_N = \mathbf{M}_N^{-1}(\hat{\boldsymbol{\beta}}_f)\mathbf{V}_{Nc}\mathbf{M}_N^{-1}(\hat{\boldsymbol{\beta}}_f)$, in which

$$\mathbf{M}_N(\boldsymbol{\beta}) = \frac{1}{N} \sum_{i=1}^N \phi_i(\boldsymbol{\beta}) \mathbf{x}_i \mathbf{x}_i^\top, \quad \mathbf{V}_{Nc} = \frac{1}{N} \sum_{i=1}^N \frac{|y_i - p_i(\hat{\boldsymbol{\beta}}_f)|^2 \mathbf{x}_i \mathbf{x}_i^\top}{N\pi_i},$$

and $\phi_i(\boldsymbol{\beta}) = p_i(\boldsymbol{\beta})\{1 - p_i(\boldsymbol{\beta})\}$ with $p_i(\boldsymbol{\beta}) = e^{\mathbf{x}_i^\top \boldsymbol{\beta}} / (1 + e^{\mathbf{x}_i^\top \boldsymbol{\beta}})$. Based on this asymptotic distribution, they derive the following two optimal subsampling probabilities

$$\pi_i^{\text{mMSE}}(\hat{\boldsymbol{\beta}}_f) = \frac{|y_i - p_i(\hat{\boldsymbol{\beta}}_f)| \|\mathbf{M}_N^{-1}(\hat{\boldsymbol{\beta}}_f) \mathbf{x}_i\|}{\sum_{j=1}^N |y_j - p_j(\hat{\boldsymbol{\beta}}_f)| \|\mathbf{M}_N^{-1}(\hat{\boldsymbol{\beta}}_f) \mathbf{x}_j\|}, \quad i = 1, \dots, N; \quad (4)$$

$$\pi_i^{\text{mVc}}(\hat{\boldsymbol{\beta}}_f) = \frac{|y_i - p_i(\hat{\boldsymbol{\beta}}_f)| \|\mathbf{x}_i\|}{\sum_{j=1}^N |y_j - p_j(\hat{\boldsymbol{\beta}}_f)| \|\mathbf{x}_j\|}, \quad i = 1, \dots, N. \quad (5)$$

Here, π_i^{mMSE} minimize the trace of \mathbf{V}_N , $\text{tr}(\mathbf{V}_N)$, and π_i^{mVc} minimize the trace of \mathbf{V}_{Nc} , $\text{tr}(\mathbf{V}_{Nc})$. These subsampling probabilities have a lot of nice properties and meaningful interpretations. More details can be found in Section 3 of Wang et al. (2017).

For ease of presentation, use the following general notation to denote subsampling probabilities

$$\pi_i^{\text{OS}}(\boldsymbol{\beta}) = \frac{|y_i - p_i(\boldsymbol{\beta})| h(\mathbf{x}_i)}{\sum_{j=1}^N |y_j - p_j(\boldsymbol{\beta})| h(\mathbf{x}_j)}, \quad i = 1, \dots, N, \quad (6)$$

where $h(\mathbf{x})$ is a univariate function of \mathbf{x} . If $h(\mathbf{x}) = \|\mathbf{x}\|$, then $\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_f)$ become $\pi_i^{\text{mVc}}(\hat{\boldsymbol{\beta}}_f)$. If $h(\mathbf{x}) = \|\mathbf{M}^{-1}\mathbf{x}\|$ where \mathbf{M} is the limit of $\mathbf{M}_N(\hat{\boldsymbol{\beta}}_f)$, then $\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_f)$ approximate $\pi_i^{\text{mMSE}}(\hat{\boldsymbol{\beta}}_f)$. If $h(\mathbf{x}) = 1$, then $\pi_i^{\text{OS}}(\boldsymbol{\beta})$ are proportional to the local case-control subsampling probabilities (Fithian and Hastie, 2014).

Note that $\pi_i^{\text{OS}}(\boldsymbol{\beta})$ depend on the unknown $\boldsymbol{\beta}$, so a pilot estimate of $\boldsymbol{\beta}$ is required to approximate them. Let $\hat{\boldsymbol{\beta}}_1$ be a pilot estimator from the pilot sample, for which we will provide more details in Section 5. The original weighted OSMAC estimator is

$$\hat{\boldsymbol{\beta}}_w = \arg \max_{\boldsymbol{\beta}} \sum_{i=1}^n \frac{y_i^* \boldsymbol{\beta}^\top \mathbf{x}_i^* - \log(1 + e^{\boldsymbol{\beta}^\top \mathbf{x}_i^*})}{\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)^*}. \quad (7)$$

In Wang et al. (2017), $\hat{\boldsymbol{\beta}}_w$ have exceptional performance because $\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)$ are able to include more informative data points in the subsample. However, the weighting scheme adopted in (7) prevent $\hat{\boldsymbol{\beta}}_w$ from being the most efficient estimator. Intuitively, a larger $\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)$ means that the data point (\mathbf{x}_i, y_i) contains more information about $\boldsymbol{\beta}$, but it has a smaller weight in the log-likelihood in (7). This reduces contributions of more informative data points in the log-likelihood function for parameter estimation.

The weighted MLE in (7) is used because $\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)$ dependent on the responses y_i 's and an un-weighted MLE is biased. If the bias can be corrected, then the resultant estimator can be more efficient in parameter estimation. This is a major goal of this paper. Interestingly, for the subsampling probabilities in (7), the bright idea proposed in Fithian and Hastie (2014) can be used to correct the bias.

3 More efficient estimator with optimal subsampling

Let $\{(\mathbf{x}_1^*, y_1^*), \dots, (\mathbf{x}_n^*, y_n^*)\}$ be a random subsample of size n taken from the full data using sampling with replacement according to the probabilities $\{\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)\}_{i=1}^N$ defined in (6). We present a more efficient estimation procedure based on subsample. Denote the more efficient estimator as $\hat{\boldsymbol{\beta}}_r$, where the subscript r stands for sampling with replacement. Remember that a pilot estimate is required, and we use $\hat{\boldsymbol{\beta}}_1$ to denote it. We will discuss how to obtain it and how to use it to improve the estimation efficiency in Section 5. Here, we focus the discussion on $\hat{\boldsymbol{\beta}}_r$ and assume that a consistent $\hat{\boldsymbol{\beta}}_1$ is available. The following procedure describes how to obtain $\hat{\boldsymbol{\beta}}_r$.

Calculate

$$\tilde{\boldsymbol{\beta}}_r = \arg \max_{\boldsymbol{\beta}} \ell_r^*(\boldsymbol{\beta}) = \arg \max_{\boldsymbol{\beta}} \sum_{i=1}^n \{\boldsymbol{\beta}^T \mathbf{x}_i^* y_i^* - \log(1 + e^{\boldsymbol{\beta}^T \mathbf{x}_i^*})\}, \quad (8)$$

and let

$$\hat{\boldsymbol{\beta}}_r = \tilde{\boldsymbol{\beta}}_r + \hat{\boldsymbol{\beta}}_1. \quad (9)$$

The naive un-weighted MLE $\tilde{\boldsymbol{\beta}}_r$ in (8) is biased, and the bias is corrected in (9) using $\hat{\boldsymbol{\beta}}_1$. We will show in the following that $\hat{\boldsymbol{\beta}}_r$ is asymptotically unbiased. This, together with the fact that $\hat{\boldsymbol{\beta}}_1$ is consistent, shows the interesting fact that $\tilde{\boldsymbol{\beta}}_r$ converges to $\mathbf{0}$ in probability as $n \rightarrow \infty$ and $N \rightarrow \infty$.

To investigate the asymptotic properties, we use $\boldsymbol{\beta}_t$ to denote the true value of $\boldsymbol{\beta}$, and summarize some regularity conditions that are required in the following.

Assumption 1. *The matrix $\mathbb{E}\{\phi(\boldsymbol{\beta}_t)h(\mathbf{x})\mathbf{x}\mathbf{x}^T\}$ is finite and positive-definite.*

Assumption 2. *The covariate \mathbf{x} and function $h(\cdot)$ satisfy that $\mathbb{E}\{\|\mathbf{x}\|^2 h^2(\mathbf{x})\} < \infty$, and $\mathbb{E}\{\|\mathbf{x}\|^2 h(\mathbf{x})\} < \infty$.*

Assumption 3. *As $n \rightarrow \infty$, $n\mathbb{E}\{h(\mathbf{x})I(\|\mathbf{x}\|^2 > n)\} \rightarrow 0$, where $I(\cdot)$ is the indicator function.*

Assumption 1 is required to establish the asymptotic normality. This is a commonly used assumption, e.g., in Fithian and Hastie (2014); Wang et al. (2017), among others. Assumptions 2 and 3 impose moment conditions on the covariate distribution and the function

$h(\mathbf{x})$. When $h(\mathbf{x}) = 1$, if $\mathbb{E}\|\mathbf{x}\|^2 < \infty$, then both the two conditions in Assumption 2 and the condition in Assumption 3 hold. Thus, the assumptions required in this paper are not stronger than those required by Fithian and Hastie (2014). When $h(\mathbf{x}) = \|\mathbf{x}\|$, by Hölder's inequality,

$$n\mathbb{E}\{h(\mathbf{x})I(\|\mathbf{x}\|^2 > n)\} \leq n(\mathbb{E}\|\mathbf{x}\|^3)^{1/3}\{\mathbb{E}I(\|\mathbf{x}\|^2 > n)\}^{2/3} = (\mathbb{E}\|\mathbf{x}\|^3)^{1/3}\{n^{3/2}\mathbb{P}(\|\mathbf{x}\|^3 > n^{3/2})\}^{2/3}.$$

Therefore, if $\mathbb{E}\|\mathbf{x}\|^3 < \infty$, Assumption 3 holds. This shows that $\mathbb{E}\|\mathbf{x}\|^4 < \infty$ implies all the three conditions required in Assumptions 2 and 3. Note that Wang et al. (2017) requires that $\mathbb{E}e^{\mathbf{v}^T\mathbf{x}} < \infty$ for any $\mathbf{v} \in \mathbb{R}^d$ in order to establish the asymptotic properties when a pilot estimate is used to approximate optimal subsampling probabilities. Thus, the required conditions in this paper are weaker than those required in Wang et al. (2017). Assumptions 1 and 2 are required in all the theorems in this paper while Assumption 3 is only required in Theorem 1.

Theorem 1. *Under assumptions 1-3, conditional on \mathcal{D}_N and $\hat{\beta}_1$, as $n \rightarrow \infty$ and $N \rightarrow \infty$,*

$$\sqrt{n}(\hat{\beta}_r - \hat{\beta}_{wf}) \longrightarrow \mathbb{N}(\mathbf{0}, \Sigma_{\beta_t}), \quad (10)$$

in distribution; furthermore, if $n/N \rightarrow 0$, then

$$\sqrt{n}(\hat{\beta}_r - \beta_t) \longrightarrow \mathbb{N}(\mathbf{0}, \Sigma_{\beta_t}) \quad (11)$$

in distribution. Here

$$\Sigma_{\beta} = \left[\frac{\mathbb{E}\{\phi(\beta)h(\mathbf{x})\mathbf{x}\mathbf{x}^T\}}{4\Phi(\beta)} \right]^{-1}, \quad \Phi(\beta) = \mathbb{E}\{\phi(\beta)h(\mathbf{x})\} \quad (12)$$

and

$$\hat{\beta}_{wf} = \arg \max_{\beta} \sum_{i=1}^N |y_i - p_i(\hat{\beta}_1)| h(\mathbf{x}_i) [y_i \mathbf{x}_i^T (\beta - \hat{\beta}_1) - \log\{1 + e^{\mathbf{x}_i^T (\beta - \hat{\beta}_1)}\}], \quad (13)$$

which satisfies that, conditional on $\hat{\beta}_1$,

$$\sqrt{N}(\hat{\beta}_{wf} - \beta_t) \longrightarrow \mathbb{N}(\mathbf{0}, \Sigma_{wf}), \quad (14)$$

in distribution, and

$$\Sigma_{wf} = [\mathbb{E}\{\phi(\beta_t)h(\mathbf{x})\mathbf{x}\mathbf{x}^T\}]^{-1} \mathbb{E}\{\phi(\beta_t)h^2(\mathbf{x})\mathbf{x}\mathbf{x}\} [\mathbb{E}\{\phi(\beta_t)h(\mathbf{x})\mathbf{x}\mathbf{x}^T\}]^{-1}. \quad (15)$$

Remark 1. *The result in (10) of Theorem 1 also implies that $\hat{\beta}_r$ is \sqrt{n} -consistent in a sense that is similar to Theorem 5 of Wang et al. (2017), namely, given \mathcal{D}_N and $\hat{\beta}_1$ in probability,*

$$\hat{\beta}_r - \hat{\beta}_{wf} = O_{P|\mathcal{D}_N, \hat{\beta}_1}(n^{-1/2}). \quad (16)$$

The $O_{P|\mathcal{D}_N, \hat{\beta}_1}(n^{-1/2})$ expression in (16) means that for any $\epsilon > 0$, there exist a δ_ϵ such that

$$\mathbb{P}\left\{\sup_n \mathbb{P}(\|\hat{\beta}_r - \hat{\beta}_{wf}\| > n^{-1/2}\delta_\epsilon | \mathcal{D}_N, \hat{\beta}_1) \leq \epsilon\right\} \rightarrow 1 \quad (17)$$

as $n, N \rightarrow \infty$. Xiong and Li (2008) showed that if a sequence is bounded in conditional probability, then it is bounded in unconditional probability, i.e., if $a_n = O_{P|\mathcal{D}_N, \hat{\beta}_1}(1)$, then $a_n = O_P(1)$. Therefore, (16) implies that $\hat{\beta}_r - \hat{\beta}_{wf} = O_P(n^{-1/2})$. Similarly, (14) implies that $\hat{\beta}_{wf} - \beta_t = O_P(N^{-1/2})$. Thus, $\hat{\beta}_r - \beta_t = O_P(n^{-1/2} + N^{-1/2}) = O_P(n^{-1/2})$, showing the \sqrt{n} -consistency of $\hat{\beta}_r$ under the unconditional distribution.

Theorem 1 shows that, asymptotically, the distribution of $\hat{\beta}_r$ given \mathcal{D}_N and $\hat{\beta}_1$ is centered around $\hat{\beta}_{wf}$ with variance-covariance matrix $n^{-1}\Sigma_{\beta_t}$, and the distribution of $\hat{\beta}_{wf}$ is centered around β_t with variance-covariance matrix $N^{-1}\Sigma_{wf}$. Thus, both $n^{-1}\Sigma_{\beta_t}$ and $N^{-1}\Sigma_{wf}$ should be considered in accessing the quality of $\hat{\beta}_r$ for estimating the true parameter β_t . However, in subsampling setting, it is expected that $n \ll N$; otherwise, the computational benefit is minimum. Thus, $n^{-1}\Sigma_{\beta_t}$ is the dominating term in quantifying the variation of $\hat{\beta}_r$. If $n/N \rightarrow 0$, then the variation of $\hat{\beta}_{wf}$ can be ignored as stated in (11).

Now we compare the estimation efficiency of $\hat{\beta}_r$ with that of the weighted estimator $\hat{\beta}_w$. With the optimal subsampling probabilities $\pi_i^{\text{OS}}(\hat{\beta}_f)$, the asymptotic variance-covariance matrix (scaled by n), \mathbf{V}_N , for the weighted estimator $\hat{\beta}_w$ has a form of $\mathbf{V}_N^{\text{OS}} = \mathbf{M}_N^{-1}(\hat{\beta}_f)\mathbf{V}_{Nc}^{\text{OS}}\mathbf{M}_N^{-1}(\hat{\beta}_f)$, in which

$$\mathbf{V}_{Nc}^{\text{OS}} = \left\{ \frac{1}{N} \sum_{i=1}^N |y_i - p_i(\hat{\beta}_f)| h(\mathbf{x}_i) \right\} \left\{ \frac{1}{N} \sum_{i=1}^N \frac{|y_i - p_i(\hat{\beta}_f)| \mathbf{x}_i \mathbf{x}_i^T}{h(\mathbf{x}_i)} \right\}.$$

Note that the full data MLE $\hat{\beta}_f$ is consistent under Assumptions 1-2. If $\mathbb{E}\{\|\mathbf{x}\|^2/h(\mathbf{x})\} < \infty$, then from Lemma 1 in the appendix and the law of large numbers, \mathbf{V}_N^{OS} converges in probability to $\mathbf{V}^{\text{OS}} = \mathbf{M}^{-1}\mathbf{V}_c^{\text{OS}}\mathbf{M}^{-1}$, where

$$\mathbf{M} = \mathbb{E}\{\phi(\beta_t)\mathbf{x}\mathbf{x}^T\} \quad \text{and} \quad \mathbf{V}_c^{\text{OS}} = 4\Phi(\beta_t)\mathbb{E}\left\{\frac{\phi(\beta_t)\mathbf{x}\mathbf{x}^T}{h(\mathbf{x})}\right\}.$$

Note that the asymptotically distribution of $\hat{\beta}_w$ given \mathcal{D}_N and $\hat{\beta}_1$ is centered around $\hat{\beta}_f$. It can be shown that under Assumptions 1-2,

$$\sqrt{N}(\hat{\beta}_f - \beta_t) \rightarrow \mathbb{N}(\mathbf{0}, \mathbf{M}^{-1}). \quad (18)$$

Thus, both $n^{-1}\mathbf{V}^{\text{OS}}$ and $N^{-1}\mathbf{M}^{-1}$ should be considered in accessing the quality of $\hat{\beta}_w$ for estimating the true parameter β_t . However, similar to the case for $\hat{\beta}_r$, $N^{-1}\mathbf{M}^{-1}$ is small compared with $n^{-1}\mathbf{V}^{\text{OS}}$ if $n \ll N$, and it is negligible if $n/N \rightarrow 0$. Therefore, the relative performance between $\hat{\beta}_r$ and $\hat{\beta}_w$ are mainly determined by the relative magnitude between \mathbf{V}^{OS} and Σ_{β_t} . We have the following result comparing \mathbf{V}^{OS} and Σ_{β_t} .

Proposition 1. *Suppose that \mathbf{M} , \mathbf{V}_c^{OS} , and Σ_{β_t} are finite and positive definite matrices. We have that*

$$\Sigma_{\beta_t} \leq \mathbf{V}^{\text{OS}}, \quad (19)$$

where the inequality is in the Loewner ordering, i.e., for positive semi-definite matrices \mathbf{A} and \mathbf{B} , $\mathbf{A} \geq \mathbf{B}$ if and only if $\mathbf{A} - \mathbf{B}$ is positive semi-definite. If $h(\mathbf{x}) = 1$, then the equality in (19) holds.

This proposition shows that $\hat{\beta}_r$ is typically more efficient than $\hat{\beta}_w$ in estimating β_t . The numerical results in Section 7 also confirm this. Note that (19) holds if $h(\mathbf{x}) = 1$. This indicates that for subsample obtained from local case-control subsampling with replacement, the weighted and unweighted estimators have the same conditional asymptotic distribution.

4 Poisson sampling

For the more efficient estimator $\hat{\beta}_r$ in Section 3 as well as the weighted estimator $\hat{\beta}_w$ in Wang et al. (2017), the subsampling procedure used is sampling with replacement, which is faster to compute than sampling without replacement for a fixed sample size. In addition, the resultant subsample are independent and identically distributed (i.i.d.) conditional on the full data. However, to implement sampling with replacement, subsampling probabilities $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$ need to be calculated all at once, and a large amount of random numbers need to be generated all at once. This may reduce the computational efficiency, and it may require a large RAM to implement the method. Furthermore, since a data point may be included multiple times in the subsample, the resultant estimator may not be the most efficient.

To enhance the computation and estimation efficiency of the subsample estimator, we consider Poisson sampling, which is also fast to compute and the resultant subsample can be independent without conditioning on the full data. Note that for subsampling with replacement, a resultant subsample is in general not independent, although it is i.i.d conditional on the full data. As another advantage with Poisson sampling, there is no need to calculate subsampling probabilities all at once, nor to generate a large amount of random numbers all at once. Furthermore, a data point cannot be included in the subsample for more than one time. A limitation of Poisson sampling is that the subsample size is always random. Due to this, we abuse the notation in this section and use n to denote the expected subsample size.

Note that $\pi_i^{\text{OS}}(\beta)$ depend on the full data through the term in the denominator, $\sum_{i=1}^N |y_i - p_i(\beta)|h(\mathbf{x}_i)$. Write $\Psi_N(\beta) = N^{-1} \sum_{i=1}^N |y_i - p_i(\beta)|h(\mathbf{x}_i)$, and denote its limit as $\Psi(\beta) = \mathbb{E}\{|y - p(\beta)|h(\mathbf{x})\}$. Note that $\Psi(\beta_t) = 2\Phi(\beta_t)$. The pilot sample can be used to obtain an estimator of $\Psi(\beta_t)$ to approximate $\Psi_N(\beta)$. Let $\hat{\Psi}_1$ be a pilot estimator of $\Psi(\beta_t)$. Here, we

focus on the Poisson sampling procedure and assume that such $\hat{\Psi}_1$ is available and consistent. We will provide more details on this in the next section.

With $\hat{\beta}_1$ and $\hat{\Psi}_1$ available, the Poisson subsampling procedure is described as the following. For $i = 1, \dots, N$, calculate $\pi_i^p = |y_i - p_i(\hat{\beta}_1)|h(\mathbf{x}_i)/(N\hat{\Psi}_1)$, generate $u_i \sim U(0, 1)$, and include $(\mathbf{x}_i, y_i, \pi_i^p)$ in the subsample if $u_i \leq n\pi_i^p$. For the obtained subsample, say $\{(\mathbf{x}_1^*, y_1^*, \pi_1^{p*}), \dots, (\mathbf{x}_{n^*}^*, y_{n^*}^*, \pi_{n^*}^{p*})\}$, calculate

$$\tilde{\beta}_p = \arg \max_{\beta} \ell_p^*(\beta) = \arg \max_{\beta} \sum_{i=1}^{n^*} (n\pi_i^{p*} \vee 1) \{\beta^T \mathbf{x}_i^* y_i^* + \log(1 + e^{\beta^T \mathbf{x}_i^*})\}, \quad (20)$$

and let $\hat{\beta}_p = \tilde{\beta}_p + \hat{\beta}_1$. Note that here the actual subsample size n^* is random.

Poisson sampling does not require to calculate π_i^p 's all at once; each π_i^p can be calculated on the go for each individual data point when scanning through the full data. Thus, one pass through the data finishes the sampling. For the estimation step, if π_i^p is large so that $n\pi_i^p > 1$, then this more informative data point will be given a larger weight, $n\pi_i^p$, in the log-likelihood in (20). The following theorem describes asymptotic properties of $\hat{\beta}_p$.

Theorem 2. *Under assumptions 1-2, conditional on \mathcal{D}_N and $\hat{\beta}_1$, as $n \rightarrow \infty$ and $N \rightarrow \infty$, if $n/N \rightarrow 0$*

$$\sqrt{n}(\hat{\beta}_p - \beta_t) \longrightarrow \mathbb{N}(0, \Sigma_{\beta_t}), \quad (21)$$

in distribution; if $n/N \rightarrow \rho \in (0, 1)$, then

$$\sqrt{n}(\hat{\beta}_p - \hat{\beta}_{wf}) \longrightarrow \mathbb{N}(0, \Sigma_{\beta_t} \Lambda_{\rho} \Sigma_{\beta_t}), \quad (22)$$

in distribution, where

$$\Lambda_{\rho} = \frac{\mathbb{E}[\phi(\beta_t)h(\mathbf{x})\{\Phi(\beta_t) - \rho\phi(\beta_t)h(\mathbf{x})\}_+ \mathbf{x}\mathbf{x}^T]}{4\Phi^2(\beta_t)}, \quad (23)$$

and $()_+$ means the positive part of the quantity, i.e., $a_+ = aI(a > 0)$.

Remark 2. *Similar to the case of Theorem 1, Theorem 2 implies that $\hat{\beta}_p$ is \sqrt{n} -consistent to β_t under the unconditional distribution.*

Theorem 2 shows that with Poisson sampling, the asymptotic variance-covariance matrices may differ for different sampling ratios n/N . In addition, comparing Theorems 1 and 2, we know that $\hat{\beta}_r$ and $\hat{\beta}_p$ have the same asymptotic distribution if $n/N \rightarrow 0$. This is intuitive because if the sampling rate n/N is small, sampling with replacement is not very different from sampling without replacement. However, if the sampling rate n/N does not converge to zero, $\hat{\beta}_r$ and $\hat{\beta}_p$ have the same asymptotic bias but different asymptotic variance-covariance matrices. The following result compares the two asymptotic variance covariance matrices.

Proposition 2. *If Σ_{β_t} is a finite and positive definite matrix and $\rho > 0$, then*

$$\Sigma_{\beta_t} \Lambda_\rho \Sigma_{\beta_t} < \Sigma_{\beta_t}, \quad (24)$$

under the Loewner ordering.

This proposition shows that Poisson sampling is more efficient than sampling with replacement if the expected subsample size is proportional to the full data size.

5 Pilot estimate and practical implementation

Since $\pi_i^{\text{OS}}(\beta)$ depend on the unknown β , a pilot estimate of β is required to approximate them. The pilot estimate can be obtained by taking a pilot sample using uniform subsampling or case-control subsampling (Wang et al., 2017). For uniform subsampling, all subsampling probabilities are equal, while for case-control subsampling, the subsampling probability for the cases ($y_i = 1$) are different from that for the controls ($y_i = 0$). Let the subsampling probabilities used to take the pilot sample be

$$\pi_{1i} = \frac{c_0(1 - y_i) + c_1 y_i}{N}, \quad (25)$$

where c_0 and c_1 are two constants that can be used to balance the numbers of 0's and 1's in the responses for the pilot subsample. If $c_0 = c_1 = 1$, then $\pi_{1i} = N^{-1}$ are the uniform subsampling probability. This choice is recommended due to its simplicity if the proportion of 1's is close to 0.5 (Wang et al., 2017). If $c_0 \neq c_1$, then π_{1i} are the case control subsampling probabilities. This choice is recommended for imbalanced full data. Often, some prior information about the marginal probability $\mathbb{P}(y = 1)$ is available. If p_{pr} is the prior marginal probability, we can choose $c_0 = \{2(1 - p_{pr})\}^{-1}$ and $c_1 = (2p_{pr})^{-1}$. The pilot estimate $\hat{\beta}_1$ can be obtained using the pilot subsample. For uniform subsampling, weighted and unweighted estimators are the same. For case-control subsampling, we use the unweighted estimators for both sampling with replacement and Poisson sampling.

In Wang et al. (2017), the pilot sample are combined with the second stage sample taken using approximated optimal subsampling probabilities to obtain a final estimator. While this does not make a difference asymptotically since n_1 is typically a small term compared with n , i.e., $n_1 = o(n)$, using the pilot sample helps to improve the finite sample performance in practical application. However, combing the raw samples may not be the most computationally efficient way of utilizing the pilot sample. Since $\hat{\beta}_1$ is already calculated, we can use it directly to improve the second stage estimator using a divide-and-conquer method. This avoids iterative calculation on the pilot sample for the second time.

For Poisson subsampling, the pilot sample can also be used to construct $\hat{\Psi}_1$ to approximate $\Psi_N(\boldsymbol{\beta})$. We use the following expression to obtain $\hat{\Psi}_1$.

$$\hat{\Psi}_1 = \frac{1}{N} \sum_{i=1}^{n_1^*} \frac{|y_i^* - p_i^*(\hat{\boldsymbol{\beta}}_1)|h(\mathbf{x}_i^*)}{n\pi_{1i}^* \wedge 1}, \quad (26)$$

where $p_i^*(\boldsymbol{\beta}) = e^{\boldsymbol{\beta}^T \mathbf{x}_i^*} / (1 + e^{\boldsymbol{\beta}^T \mathbf{x}_i^*})$. It can be verified that $\hat{\Psi}_1$ defined in (26) converges in probability to $\Psi(\boldsymbol{\beta}_t)$.

Another practical issue one has to consider is whether the full data can be loaded into available RAM or not. If so, subsampling probabilities can be calculated in RAM and subsampling with replacement can be implemented directly. Otherwise, special considerations have to be given in practical implementation. If the full data is larger than available RAM while subsampling probabilities $\{\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)\}_{i=1}^N$ can still be loaded in available RAM, one can calculate $\{\pi_i^{\text{OS}}(\hat{\boldsymbol{\beta}}_1)\}_{i=1}^N$ by scanning the data from hard drive line-by-line or block-by-block, generate row indexes for a subsample, and then scan the data line-by-line or block-by-block to take the subsample. To be specific, one can draw a subsample, say $\{idx_1, \dots, idx_n\}$, from $\{1, \dots, N\}$, sort the indexes to have $\{idx_{(1)}, \dots, idx_{(n)}\}$, and then use the following algorithm to scan the data line-by-line or block-by-block in order to obtain the subsample.

Algorithm 1 Obtain subsample with given indexes by scanning through the full data

Input: data file, subsample indexes $\{idx_{(1)}, \dots, idx_{(n)}\}$.

$i \leftarrow 1$

$j \leftarrow 1$

while $i \leq N$ **and** $j \leq n$ **do**

 readline(data file)

if $i == idx_{(j)}$ **then**

 include the i th data point into the subsample

while $i == idx_{(j)}$ **do**

$j \leftarrow j + 1$

end while

end if

$i \leftarrow i + 1$

end while

Remark 3. Clearly, Algorithm 1 takes no more than linear time to run. We assume that a generic function `readline()` reads a single line (or multiple lines) from the data file and stop at the beginning of the next line (or next block) in the data file. No calculation is performed

on a data line if it is not included in the subsample. Such functionality are provided by most programming languages. For example, Julia (Bezanson et al., 2017) and Python (van Rossum, 1995) has a function `readline()` that read a file line-by-line; R (R Core Team, 2017) has a function `readLines()` that read one or multiple lines; C (Kernighan and Ritchie, 1988) and C++ (Stroustrup, 1986) has a function `getline()` to read one line at a time.

Taking into account all aforementioned issues in this section, including how to obtain the pilot estimates, how to combine it with the second stage estimates, as well as how to process data file line-by-line, we summarize practical implementation procedures in Algorithm 2 and Algorithm 3, where Algorithm 2 is for sampling with replacement and Algorithm 3 is for Poisson sampling. The algorithms are presented for the scenario that data volume is larger than the size the the RAM. When the full data can be loaded into the RAM, the procedure of sampling can be done directly in the RAM and the method is easier to implement for that scenario.

Remark 4. In Algorithm 2 and Algorithm 3, if $n_1 = o(n)$, then the result for $\hat{\beta}_r$ in Theorem 1 hold for $\check{\beta}_r$ and the result for $\hat{\beta}_p$ in Theorem 2 hold for $\check{\beta}_p$ as well. This is because $\{\check{\ell}_r^{*1}(\check{\beta}_1) + \check{\ell}_r^*(\check{\beta}_r)\}^{-1}\check{\ell}_r^{*1}(\check{\beta}_1)\sqrt{n}(\hat{\beta}_1 - \beta_t) = O_p(\sqrt{n_1}/\sqrt{n}) = o_P(1)$ and $\{\check{\ell}_r^{*1}(\check{\beta}_1) + \check{\ell}_r^*(\check{\beta}_r)\}^{-1}\check{\ell}_r(\check{\beta}_r) \rightarrow 1$ in probability. The reason for $\check{\beta}_p$ is similar.

Remark 5. In Algorithm 2 and Algorithm 3, to combine the two stage estimates using the second derivative of the likelihood, the inconsistent estimators $\check{\beta}_1$, and $\check{\beta}_r$ or $\check{\beta}_p$ should be used, because their limits correspond to the terms in the asymptotic variance-covariance matrices of the more efficient estimators. This is an advantage of the proposed estimators for implementation using existing software that fit logistic regression. One can use the inverse of the estimated variance-covariance matrix from the software output to replace the second derivative of the likelihood.

Remark 6. The variance-covariance estimators $\hat{V}(\check{\beta}_r)$ in (31) and $\hat{V}(\check{\beta}_p)$ in (35) can be replaced the following simplified estimators,

$$\hat{V}_s(\check{\beta}_p) = \{\check{\ell}_p^{*1}(\check{\beta}_1) + \check{\ell}_p^*(\check{\beta}_p)\}^{-1} \quad \text{and} \quad \hat{V}_s(\check{\beta}_r) = \{\check{\ell}_r^{*1}(\check{\beta}_1) + \check{\ell}_r^*(\check{\beta}_r)\}^{-1}, \quad (27)$$

respectively. If the subsampling ratio n/N is much smaller than one, then $\hat{V}_s(\check{\beta}_r)$ and $\hat{V}_s(\check{\beta}_p)$ perform very similarly $\hat{V}(\check{\beta}_r)$ and $\hat{V}(\check{\beta}_p)$, respectively.

Remark 7. The time complexity of Algorithm 2 is the same as that of Algorithm 2 in Wang et al. (2017). The major computing time is to calculate $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$ in Step 2, but it does not require iterative calculations on the full data. Once $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$ are available, it takes

$O(N)$ time to obtain the subsample using Algorithm 1 and the calculations of $\hat{\beta}_{uw}$ and $\check{\beta}_{uw}$ are fast because they are done on the subsamples only. To calculate $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$, the required time varies. For π_i^{mVc} , the required time is $O(Nd)$; for π_i^{mMSE} , the required time is $O(Nd^2)$. Thus, the time complexity of Algorithm 2 with π_i^{mVc} is $O(Nd)$ and the time complexity with π_i^{mMSE} is $O(Nd^2)$, if the sampling ratio n/N is much smaller than one.

Algorithm 2 More efficient estimation based on subsampling with replacement

Step 1: obtain the pilot $\hat{\beta}_1$

- (1) Take pilot subsample indexes of size n_1 from $\{1, \dots, N\}$ using sampling with replacement according to subsampling probabilities π_{1i} in (25), and use Algorithm 1 to obtain a pilot subsample, say, $(\mathbf{x}_i^{*1}, y_i^{*1})$, $i = 1, \dots, n_1$.

- (2) Calculate

$$\tilde{\beta}_1 = \arg \max_{\beta} \ell_r^{*1}(\beta) = \arg \max_{\beta} \sum_{i=1}^{n_1} \{\beta^T \mathbf{x}_i^{*1} y_i^{*1} - \log(1 + e^{\beta^T \mathbf{x}_i^{*1}})\}, \quad (28)$$

and let $\hat{\beta}_1 = \tilde{\beta}_1 + \mathbf{b}$, where $\mathbf{b} = \{\log(c_0/c_1), 0, \dots, 0\}^T$.

Step 2: obtain the more efficient estimator $\hat{\beta}_r$

- (1) Calculate $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$ defined in equation (6); take subsample indexes of size n from $\{1, \dots, N\}$ according to sampling probabilities $\{\pi_i^{\text{OS}}(\hat{\beta}_1)\}_{i=1}^N$ using sampling with replacement; and use Algorithm 1 to obtain a subsample, (\mathbf{x}_i^*, y_i^*) , $i = 1, \dots, n$.

- (2) Calculate

$$\tilde{\beta}_r = \arg \max_{\beta} \ell_r^*(\beta) = \arg \max_{\beta} \sum_{i=1}^n \{\beta^T \mathbf{x}_i^* y_i^* - \log(1 + e^{\beta^T \mathbf{x}_i^*})\}, \quad (29)$$

and let $\hat{\beta}_r = \tilde{\beta}_r + \hat{\beta}_1$.

Step 3: combine the two estimators $\hat{\beta}_1$ and $\hat{\beta}_r$

Calculate

$$\check{\beta}_r = \{\ddot{\ell}_r^{*1}(\tilde{\beta}_1) + \ddot{\ell}_r^*(\tilde{\beta}_r)\}^{-1} \{\ddot{\ell}_r^{*1}(\tilde{\beta}_1)\hat{\beta}_1 + \ddot{\ell}_r^*(\tilde{\beta}_r)\hat{\beta}_r\} \quad (30)$$

where $\ddot{\ell}_r^{*1}(\tilde{\beta}_1) = \sum_{i=1}^{n_1} \phi^{*1}(\tilde{\beta}_1) \mathbf{x}_i^{*1} (\mathbf{x}_i^{*1})^T$ and $\ddot{\ell}_r^*(\tilde{\beta}_r) = \sum_{i=1}^n \phi^*(\tilde{\beta}_r) \mathbf{x}_i^* (\mathbf{x}_i^*)^T$.

The variance-covariance matrix of $\check{\beta}_r$ can be estimated by

$$\hat{V}(\check{\beta}_r) = \{\ddot{\ell}_r^{*1}(\tilde{\beta}_1) + \ddot{\ell}_r^*(\tilde{\beta}_r)\}^{-1} \left[\sum_{i=1}^{n_1} \{\psi_i^{*1}(\tilde{\beta}_1)\}^2 \mathbf{x}_i^{*1} (\mathbf{x}_i^{*1})^T \times \right. \\ \left. + \sum_{i=1}^n \{\psi_i^*(\tilde{\beta}_r)\}^2 \mathbf{x}_i^* (\mathbf{x}_i^*)^T \right] \{\ddot{\ell}_r^{*1}(\tilde{\beta}_1) + \ddot{\ell}_r^*(\tilde{\beta}_r)\}^{-1}, \quad (31)$$

where $\psi_i^*(\beta) = y_i^* - p_i^*(\beta)$.

Algorithm 3 More efficient estimation based on Poisson sampling

Step 1: obtain the pilots $\hat{\beta}_1$ and $\hat{\Psi}_1$

- (1) For $i = 1, \dots, N$, calculate $\pi_{1i} = \frac{c_0(1-y_i)+c_1y_i}{N}$, generate $u_{1i} \sim U(0, 1)$, and add $(\mathbf{x}_i, y_i, \pi_{1i})$ in the subsample if $u_{1i} \leq n_1\pi_{1i}$.
- (2) For the obtained subsample, say $(\mathbf{x}_i^{*1}, y_i^{*1}, \pi_{1i}^{*1})$, $i = 1, \dots, n_1^*$, calculate

$$\tilde{\beta}_1 = \arg \max_{\beta} \ell_p^{*1}(\beta) = \arg \max_{\beta} \sum_{i=1}^{n_1^*} (n\pi_{1i}^{*1} \vee 1) \{ \beta^T \mathbf{x}_i^{*1} y_i^{*1} + \log(1 + e^{\beta^T \mathbf{x}_i^{*1}}) \}, \quad (32)$$

let $\hat{\beta}_1 = \tilde{\beta}_1 + \mathbf{b}$, and then calculate $\hat{\Psi}_1$ in equation (26).

Step 2: obtain the more efficient estimator $\hat{\beta}_p$

- (1) For $i = 1, \dots, N$, calculate $\pi_i^p = \frac{|y_i - p_i(\hat{\beta}_1)|h(\mathbf{x}_i)}{N\hat{\Psi}_1}$, generate $u_i \sim U(0, 1)$, and if $u_i \leq n\pi_i^p$ add $(\mathbf{x}_i, y_i, \pi_i^p)$ in the subsample.
- (2) For the obtained subsample, say $\{(\mathbf{x}_1^*, y_1^*, \pi_1^{p*}), \dots, (\mathbf{x}_{n^*}^*, y_{n^*}^*, \pi_{n^*}^{p*})\}$, calculate

$$\tilde{\beta}_p = \arg \max_{\beta} \ell_p^*(\beta) = \arg \max_{\beta} \sum_{i=1}^{n^*} (n\pi_i^{p*} \vee 1) \{ \beta^T \mathbf{x}_i^* y_i^* + \log(1 + e^{\beta^T \mathbf{x}_i^*}) \}, \quad (33)$$

and let $\hat{\beta}_p = \tilde{\beta}_p + \hat{\beta}_1$.

Step 3: combine the two estimators $\hat{\beta}_1$ and $\hat{\beta}_p$

Calculate

$$\check{\beta}_p = \{ \check{\ell}_p^{*1}(\tilde{\beta}_1) + \check{\ell}_p^*(\tilde{\beta}_p) \}^{-1} \{ \check{\ell}_p^{*1}(\tilde{\beta}_1) \hat{\beta}_1 + \check{\ell}_p^*(\tilde{\beta}_p) \hat{\beta}_p \} \quad (34)$$

where $\check{\ell}_p^{*1}(\tilde{\beta}_1) = \sum_{i=1}^{n_1^*} \phi^{*1}(\tilde{\beta}_1) \mathbf{x}_i^{*1} (\mathbf{x}_i^{*1})^T$ and $\check{\ell}_p^*(\tilde{\beta}_p) = \sum_{i=1}^{n^*} \phi^*(\tilde{\beta}_p) \mathbf{x}_i^* (\mathbf{x}_i^*)^T$.

The variance-covariance matrix of $\check{\beta}_p$ can be estimated by

$$\hat{V}(\check{\beta}_p) = \{ \check{\ell}_p^{*1}(\tilde{\beta}_1) + \check{\ell}_p^*(\tilde{\beta}_p) \}^{-1} \left[\sum_{i=1}^{n_1^*} \{ \psi_i^{*1}(\tilde{\beta}_1) \}^2 \mathbf{x}_i^{*1} (\mathbf{x}_i^{*1})^T \times \right. \\ \left. + \sum_{i=1}^{n^*} \{ \psi_i^*(\tilde{\beta}_p) \}^2 \mathbf{x}_i^* (\mathbf{x}_i^*)^T \right] \{ \check{\ell}_p^{*1}(\tilde{\beta}_1) + \check{\ell}_p^*(\tilde{\beta}_p) \}^{-1}. \quad (35)$$

6 Unconditional distribution

Asymptotic distributional results in Sections 3 and 4, as well as in Wang et al. (2017), are about conditional distributions, i.e., they are about conditional distributions of subsample-based estimators given the full data and the pilot estimate. We investigate the unconditional distribution of $\hat{\beta}_p$ in this section.

Theorem 3. *Assume that Assumptions 1 and 2 hold. If the pilot estimators are obtained with subsampling probabilities in (25) and $\mathbb{E}\{h^3(\mathbf{x})\|\mathbf{x}\|^3\}$, $\mathbb{E}\{h^3(\mathbf{x})\|\mathbf{x}\|^2\}$, $\mathbb{E}\{h(\mathbf{x})\|\mathbf{x}\|^3\}$, and $\mathbb{E}\{h^2(\mathbf{x})\}$ are finite, or if $\hat{\beta}_1$ and $\hat{\Psi}_1$ are independent of the data, then the following results hold.*

As $n \rightarrow \infty$ and $N \rightarrow \infty$, if $n/N \rightarrow 0$

$$\sqrt{n}(\hat{\beta}_p - \beta_t) \longrightarrow \mathbb{N}(0, \Sigma_{\beta_t}), \quad (36)$$

in distribution; if $n/N \rightarrow \rho \in (0, 1)$, then

$$\sqrt{n}(\hat{\beta}_p - \beta_t) \longrightarrow \mathbb{N}(0, \Sigma_{\beta_t} \Lambda_u \Sigma_{\beta_t}), \quad (37)$$

in distribution, where

$$\Lambda_u = \frac{\mathbb{E}[\phi(\beta_t)\{\rho\phi(\beta_t)h(\mathbf{x}) \vee \Phi(\beta_t)\}h(\mathbf{x})\mathbf{x}\mathbf{x}^T]}{4\Phi^2(\beta_t)}. \quad (38)$$

Remark 8. *If the pilot estimators $\hat{\beta}_1$ and $\hat{\Psi}_1$ are obtained through the data, stronger moment conditions are required. Note that $h(\mathbf{x})$ is often a function of the norm of \mathbf{x} , such as in π_i^{mVc} , π_i^{mMSE} , and the local case control subsampling. In general, if $h(\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|^a$ for some constant matrix \mathbf{A} and constant $a \geq 0$, then the four additional moment conditions reduce to one requirement of $\mathbb{E}\{h^3(\mathbf{x})\|\mathbf{x}\|^3\} < \infty$.*

Fithian and Hastie (2014) obtained unconditional distribution of local case-control estimator by assuming that the pilot estimate is independent of the data. Our Theorem 3 includes this scenario, and the required assumptions are the same as those required in Fithian and Hastie (2014). In practice, a consistent pilot estimator that is independent of the data may not be available and a pilot subsample from the full data is required to construct it. For this scenario, a pilot estimator is dependent on the data, and we need a stronger moment condition to establish the asymptotic normality. For local case-control subsampling, $h(\mathbf{x}) = 1$, so the additional moment requirement is that $\mathbb{E}(\|\mathbf{x}\|^3) < \infty$.

From Theorems 2 and 3, the conditional asymptotic distribution and unconditional asymptotic distribution of $\hat{\beta}_p$ are the same if $n/N \rightarrow 0$. This is intuitive, because if the

sampling rate n/N is small, the variation of $\hat{\beta}_p$ due to the variation of the full data is small compared with the variation due to the variation of the subsampling.

However, if the sampling rate n/N does not converge to zero, then the conditional asymptotic distribution and unconditional asymptotic distribution of $\hat{\beta}_p$ are quite different. First, we notice that under the unconditional distribution, $\hat{\beta}_p$ is asymptotically unbiased to β_t , while under the conditional distribution, $\hat{\beta}_p$ is asymptotically biased with the bias being $\hat{\beta}_{wf} - \beta_t = O_P(N^{-1/2})$. Second, since the variation of $\hat{\beta}_p$ due to the variation of the full data is not negligible, we expect that the asymptotic variance covariance matrices for the unconditional distribution is larger than that for the conditional distribution. Indeed this is true, and we present it in the following proposition.

Proposition 3. *If Σ_{β_t} is a finite and positive definite matrix and $\rho > 0$, then*

$$\Sigma_{\beta_t} \Lambda_u \Sigma_{\beta_t} \geq \Sigma_{\beta_t} > \Sigma_{\beta_t} \Lambda_\rho \Sigma_{\beta_t}, \quad (39)$$

under the Loewner ordering. Furthermore, if $\mathbb{P}\{\rho\phi(\beta_t)h(\mathbf{x}) > \Phi(\beta_t)\} > 0$, then the “ \geq ” sign in (39) can be replaced by “ $>$ ”, the strict great sign.

7 Numerical evaluations

We evaluate the performance of the more efficient estimators in terms of both estimation efficiency and computational efficiency in this section.

7.1 Estimation efficiency

In this section, we use numerical experiments based on simulated and real data sets to evaluate the estimators proposed in this paper. For simulation, to compare with the original OSMAC estimator, we use exactly the same setup used in Section 5.1 of Wang et al. (2017). Specifically, the full data sample size $N = 10,000$ and the true value of β , β_t , is a 7×1 vector of 0.5. The following 6 distributions of \mathbf{x} are considered: multivariate normal distribution with mean zero (mzNormal), multivariate normal distribution with nonzero mean (nzNormal), multivariate normal distribution with mean zero and unequal variances (ueNormal), mixture of two multivariate normal distributions with different means (mixNormal), multivariate t distribution with degrees of freedom 3 (T_3), and exponential distribution (EXP). Detailed explanations of these distributions can be found in Section 5.1 of Wang et al. (2017).

To evaluate the estimation performance of the new estimators compared with the original weighted OSMAC estimator, $\check{\beta}_w$, we define the estimation efficiency of $\check{\beta}_{\text{new}}$ relative to $\check{\beta}_w$

as

$$\text{Relative Efficiency} = \frac{\text{MSE}(\check{\beta}_w)}{\text{MSE}(\check{\beta}_{\text{new}})}, \quad (40)$$

where $\check{\beta}_{\text{new}} = \check{\beta}_r$ for the subsampling with replacement estimator described in Algorithm 2 and $\check{\beta}_{\text{new}} = \check{\beta}_p$ for Poisson subsampling estimator described in Algorithm 3. We calculate empirical MSEs from $S = 1000$ subsamples using $\text{MSE}(\check{\beta}) = S^{-1} \sum_{s=1}^S \|\check{\beta}^{(s)} - \beta_t\|^2$, where $\check{\beta}^{(s)}$ is the estimate from the s th subsample. We fixed the first step sample size $n_1 = 200$ and choose n to be 100, 200, 400, 600, 800, and 1000. This is the same setup used in Wang et al. (2017).

Figure 1 presents the relative efficiency of $\check{\beta}_r$ and $\check{\beta}_p$ based on two different choices of π_i^{OS} : π_i^{mMSE} and π_i^{mVc} . It is seen that in general $\check{\beta}_r$ and $\check{\beta}_p$ are more efficient than $\check{\beta}_w$. Among the six cases, the only case for $\check{\beta}_w$ to be more efficient is when \mathbf{x} has a T_3 distribution and π^{mVc} is used, but the difference is not very significant. For all other cases, $\check{\beta}_r$ and $\check{\beta}_p$ are more efficient. For example, when \mathbf{x} has a the nzNormal distribution, $\check{\beta}_p$ can be 250% as efficient as $\check{\beta}_w$ if π^{mMSE} is used. Between $\check{\beta}_r$ and $\check{\beta}_p$, $\check{\beta}_p$ is more efficient than $\check{\beta}_r$ for all cases. We also calculate the empirical unconditional MSE by generating the full data in each repetition of the simulation. The results are similar and thus are omitted.

To assess the performance of $\hat{\mathbb{V}}(\check{\beta}_r)$ in (31) and $\hat{\mathbb{V}}(\check{\beta}_p)$ in (35), we use $\text{tr}\{\hat{\mathbb{V}}(\check{\beta}_r)\}$ and $\text{tr}\{\hat{\mathbb{V}}(\check{\beta}_p)\}$ to estimate the MSEs of $\check{\beta}_r$ and $\check{\beta}_p$, and compare the average estimated MSEs with the unconditional empirical MSEs. We focus on the unconditional MSE because conditional inference may be appropriate only if $n/N \rightarrow 0$. Figure 2 presents the results for using π^{mVc} . Results for using π^{mMSE} are similar and are omitted for clearer presentation. It is seen that the estimated MSEs are very close to the empirical MSEs, except for the case of nzNormal covariate for subsampling with replacement. For this case, the responses are imbalanced with about 95% being 1's. For this scenario, the variance-covariance estimator for $\check{\beta}_w$ proposed in Wang et al. (2017) also has a similar problem of underestimation. For Poisson sampling, the problem of underestimation from $\hat{\mathbb{V}}(\check{\beta}_p)$ is not significant.




Figure 1: Relative efficiency for different second step subsample size n with the first step subsample size being fixed at $n_1 = 200$. A relative efficiency larger than one means the associate method is more efficient than the original OSMAC estimator $\hat{\beta}_w$.




Figure 2: MSE and estimated MSE, $\text{tr}\{\hat{V}(\tilde{\beta})\}$, for different second step subsample size n with the first step subsample size being fixed at $n_1 = 200$.

We also apply the more efficient estimation methods to a supersymmetric (SUSY) bench-

mark data set (Baldi et al., 2014) available from the Machine Learning Repository (Lichman, 2013). The data contains a binary response variable indicating whether a process produce new supersymmetric particles or not and 18 covariates that are kinematic features about the process. The full sample size is $N = 5,000,000$ and the data file is about 2.4 gigabytes. About 54.24% of the responses in the full data are from the background process. We use the more efficient estimation methods with subsample size n to estimate parameters in logistic regression.

Figures 3 gives the relative efficiency of $\check{\beta}_r$ and $\check{\beta}_p$ to $\check{\beta}_w$ for both π^{mVc} and π^{mMSE} . It is seen that when π^{mMSE} is used, $\check{\beta}_r$ and $\check{\beta}_p$ always outperform $\check{\beta}_w$. When π^{mVc} is used, $\check{\beta}_r$ and $\check{\beta}_p$ may not be as efficient as $\check{\beta}_w$, but they become more efficient when the second stage sample size n gets larger. It is also seen that $\check{\beta}_p$ dominates $\check{\beta}_r$ and π^{mMSE} dominates π^{mVc} in estimation efficiency.




Figure 3: Relative efficiency for the SUSY data set with $n_1 = 200$ and different second step subsample size n . The gray horizontal dashed line is the reference line when relative efficiency is one.

7.2 Computational efficiency

We consider the computational efficiency of the more efficient estimation methods in this section. Note that they have the same order of computational time complicity, so they should have similar computational efficiency. For Poisson sampling, there is no need to calculate

$\{\pi_i^p\}_{i=1}^N$ all at once and random numbers can be generated on the go, so it requires less RAM and may require less CPU times as well. To confirm this, we record the computing time of implementing each of them for the case when \mathbf{x} is `mzNormal`. All methods are implemented in the R programming language (R Core Team, 2017), and computations are carried out on a desktop running Ubuntu Linux 16.04 with an Intel I7 processor and 16GB RAM. Only one logical CPU is used for the calculation. We set the value of d to $d = 50$, the values of N to be $N = 10^4, 10^5, 10^6$ and 10^7 , and the subsample sizes to be $n_1 = 200$ and $n = 1000$.

Table 1 gives the required CPU times (in seconds) to obtain $\check{\beta}_w$, $\check{\beta}_r$, and $\check{\beta}_p$, using π^{mVc} and π^{mMSE} . The computing times for using the full data (Full) are also given for comparisons. It is seen that $\check{\beta}_r$ and $\check{\beta}_p$ are a little more computationally efficient than $\check{\beta}_w$ but the advantages are not very significant. Note that these times are obtained when all the calculations are done in the RAM, and only the CPU times for implementing each method are counted while the time to generate the data is not counted.

Table 1: CPU seconds when the full data are generate and kept in the RAM. Here $n_1 = 200$, $n = 1000$, and the full data size N varies; the covariates are from a $d = 50$ dimensional multivariate normal distribution.

| Method | N | | | |
|-------------------------|--------|--------|--------|--------|
| | 10^4 | 10^5 | 10^6 | 10^7 |
| mVc, $\check{\beta}_w$ | 0.14 | 0.13 | 0.45 | 5.24 |
| mVc, $\check{\beta}_r$ | 0.08 | 0.11 | 0.41 | 3.71 |
| mVc, $\check{\beta}_p$ | 0.08 | 0.11 | 0.43 | 3.88 |
| mMSE, $\check{\beta}_w$ | 0.13 | 0.32 | 3.31 | 35.15 |
| mMSE, $\check{\beta}_r$ | 0.12 | 0.31 | 3.29 | 34.98 |
| mMSE, $\check{\beta}_p$ | 0.12 | 0.31 | 3.29 | 35.06 |
| Full | 0.15 | 1.62 | 15.05 | 247.89 |

For big data problem, it is common that the full data are larger than the size of the available RAM, and full data can not be loaded into the RAM. For this scenario, one has to load the data into RAM line-by-line or block-by-block. Note that communication between CPU and hard drive is much slower than communication between CPU and RAM. Thus, this will dramatically increase the computing time. To mimic this situation, we store the full data on hard drive and use `readlines()` function to process data 1000 rows each time. We also use a smaller computer with 8GB RAM to implement the method. For the case when $N = 10^7$, the full data is about 9.1GB which is larger than the available RAM.

The computing times when data are scanned from hard drive are reported Table 2. Here

the computing times can be over thousand times longer than those when data are loaded into RAM. Note that these computing times can be reduced dramatically if we use some other programming language like C++ (Stroustrup, 1986) or Julia (Bezanson et al., 2017). However, for fair comparisons, we use the same programming language R here. Furthermore, our main purpose here is to demonstrate the computational advantage of subsampling so the real focus is on the relative performance among different methods. From Table 2, it is seen that using π^{mMSE} does not cost much more time than using π^{mVc} . The reason for this observation is that the major computing time is spent in data processing and the computing times used in calculating the subsampling probabilities are short. We also notice that Poisson sampling is more computational efficient than subsampling with replacement since it calculates subsampling probabilities and generates random numbers on the go and requires one time less to scan the full data. Poisson subsampling only used about 2% of the time required by implementing the full data approach.

Table 2: CPU seconds when the full data are scanned from hard drive. Here $n_1 = 200$, $n = 1000$, and the full data size N varies; the covariates are from a $d = 50$ dimensional multivariate normal distribution.

| Method | N | | | |
|-------------------------|--------|---------|----------|-----------|
| | 10^4 | 10^5 | 10^6 | 10^7 |
| mVc, $\check{\beta}_w$ | 4.26 | 41.60 | 441.46 | 4374.94 |
| mVc, $\check{\beta}_r$ | 4.13 | 41.42 | 413.09 | 4384.99 |
| mVc, $\check{\beta}_p$ | 2.77 | 27.58 | 272.32 | 2699.13 |
| mMSE, $\check{\beta}_w$ | 4.43 | 41.75 | 434.96 | 4393.38 |
| mMSE, $\check{\beta}_r$ | 4.10 | 41.83 | 417.55 | 4369.04 |
| mMSE, $\check{\beta}_p$ | 2.88 | 27.93 | 273.24 | 2719.51 |
| Full | 139.46 | 1411.78 | 14829.63 | 138134.69 |

8 Summary

In this paper, we proposed a new unweighted estimator for logistic regression based on OSMAC subsample. We derived conditional asymptotic distribution of the new estimator which has a smaller variance-covariance matrix compared with the weighted estimator.

We also investigate the asymptotic properties if Poisson sampling is used, and showed that the resultant estimator has the same conditional asymptotic distribution if the subsampling rate converges to zero. However, if the subsampling rate converges to a positive constant, the

estimator based on Poisson sampling has a smaller variance-covariance matrix.

Furthermore, we also derive the unconditional asymptotic distribution for the proposed estimator based on Poisson sampling. Interestingly, if the subsampling rate converges to zero, the unconditional asymptotic distribution is the same as the conditional asymptotic distribution, indicating that the variation of the full data can be ignored. If the subsampling rate does not converge to zero, the unconditional asymptotic distribution has a larger variance-covariance matrix. Our results also include the case-control sampling method. With a stronger moment condition that the third moment of the covariate is finite, we do not require the pilot estimate to be independent of the data.

A Proofs and technical details

In this appendix, we provide proofs for the results in the paper. Technical details related to sampling with replacement in Section 3 are presented in Section A.1; technical details related to Poisson sampling in Section 4 are presented in Section A.2; and technical details related to unconditional results in Section 6 are presented in Section A.3.

A.1 Proofs for subsampling with replacement

In this section we prove the results in Section 3. The proofs require a series of Lemmas, which are presented below and will be proved later in this section.

For easy of presentation, we use notation λ to denote the log-likelihood shifted by $\hat{\beta}_1$. For the subsample, $\lambda_r^*(\beta) = \ell_r^*(\beta - \hat{\beta}_1)$. Denote the first and second derivatives of $\lambda_r^*(\beta)$ as $\dot{\lambda}_r^*(\beta) = \partial \lambda_r^*(\beta) / \partial \beta$ and $\ddot{\lambda}_r^*(\beta) = \partial^2 \lambda_r^*(\beta) / (\partial \beta \partial \beta^T)$.

Note that from Xiong and Li (2008), the fact that a sequence converges to 0 in conditional probability is equivalent to the fact that it converges to 0 in unconditional probability. Thus, in the following, we will not use $o_P(1)$ to denote a sequence converging to 0 in probability without stating whether the underlying probability measure is conditional or unconditional.

Lemma 1. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be i.i.d. random vector with the same distribution of \mathbf{v} . Let g_1 be a bounded function that may depend on n and other random variables, and g_2 be a fixed function that does not depend on n . If $g_1(\mathbf{v}_i) = o_P(1)$ for each i as $n \rightarrow \infty$, and $\mathbb{E}|g_2(\mathbf{v})| < \infty$, then*

$$\frac{1}{n} \sum_{i=1}^n g_1(\mathbf{v}_i) g_2(\mathbf{v}_i) = o_P(1).$$

Lemma 2. Let $\eta_i = |\psi_i(\hat{\boldsymbol{\beta}}_1)|\psi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)h(\mathbf{x}_i)\mathbf{x}_i$, where $\psi_i(\boldsymbol{\beta}) = y_i - p_i(\boldsymbol{\beta})$. Under Assumptions 1 and 2, conditional on $\hat{\boldsymbol{\beta}}_1$,

$$\sqrt{N}(\hat{\boldsymbol{\beta}}_{wf} - \boldsymbol{\beta}_t) = \frac{\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}}{2\mathbb{E}\{\phi(\boldsymbol{\beta}_t)h(\mathbf{x})\}} \frac{1}{\sqrt{N}} \sum_{i=1}^N \eta_i + o_P(1), \quad (\text{A.1})$$

which, as $N \rightarrow \infty$, converges in distribution to the a normal distribution with mean $\mathbf{0}$ and variance-covariance matrix $[\mathbb{E}\{\phi(\boldsymbol{\beta}_t)h(\mathbf{x})\mathbf{x}\mathbf{x}^\top\}]^{-1}\mathbb{E}\{\phi(\boldsymbol{\beta}_t)h^2(\mathbf{x})\mathbf{x}\mathbf{x}\}[\mathbb{E}\{\phi(\boldsymbol{\beta}_t)h(\mathbf{x})\mathbf{x}\mathbf{x}^\top\}]^{-1}$.

Lemma 3. Let

$$\dot{\lambda}_r^*(\boldsymbol{\beta}_t) = \sum_{i=1}^n \{y_i^* - p_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\}\mathbf{x}_i^*. \quad (\text{A.2})$$

Under Assumptions 1 and 2, conditional on \mathcal{D}_N and $\hat{\boldsymbol{\beta}}_1$ which is assumed to be a consistent estimator of $\boldsymbol{\beta}_t$, as $n \rightarrow \infty$ and $N \rightarrow \infty$,

$$\frac{\dot{\lambda}_r^*(\boldsymbol{\beta}_t)}{\sqrt{n}} - \frac{\sqrt{n} \sum_{i=1}^N \eta_i}{N\Psi_N(\hat{\boldsymbol{\beta}}_1)} \longrightarrow \mathbb{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}), \quad (\text{A.3})$$

in distribution.

Lemma 4. Under Assumptions 1-3, as $n, N \rightarrow \infty$, for any $\mathbf{s}_n \rightarrow 0$,

$$\frac{1}{n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i^*\|^2 - \sum_{i=1}^N \pi_i(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \|\mathbf{x}_i\|^2 = o_P(1). \quad (\text{A.4})$$

Proof of Theorem 1. The estimator $\hat{\boldsymbol{\beta}}_r$ is the maximizer of

$$\lambda_r^*(\boldsymbol{\beta}) = \sum_{i=1}^n \left[(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^\top \mathbf{x}_i^* y_i^* - \log \{1 + e^{(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^\top \mathbf{x}_i^*}\} \right], \quad (\text{A.5})$$

so $\sqrt{n}(\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_t)$ is the maximizer of $\gamma(\mathbf{s}) = \lambda_r^*(\boldsymbol{\beta}_t + \mathbf{s}/\sqrt{n}) - \lambda_r^*(\boldsymbol{\beta}_t)$. By Taylor's expansion,

$$\gamma(\mathbf{s}) = \frac{1}{\sqrt{n}} \mathbf{s}^\top \dot{\lambda}_r^*(\boldsymbol{\beta}_t) + \frac{1}{2n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \dot{\mathbf{s}}/\sqrt{n}) (\mathbf{s}^\top \mathbf{x}_i^*)^2, \quad (\text{A.6})$$

where $\phi_i^*(\boldsymbol{\beta}) = p_i^*(\boldsymbol{\beta})\{1 - p_i^*(\boldsymbol{\beta})\}$, and $\dot{\mathbf{s}}$ lies between $\mathbf{0}$ and \mathbf{s} .

From Lemma 4,

$$\frac{1}{n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \dot{\mathbf{s}}/\sqrt{n}) \mathbf{x}_i^* (\mathbf{x}_i^*)^\top - \sum_{i=1}^N \pi(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i \mathbf{x}_i^\top = o_P(1). \quad (\text{A.7})$$

From Lemma 1 and the law of large numbers,

$$\sum_{i=1}^N \pi(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i \mathbf{x}_i^\top = \frac{\frac{1}{N} \sum_{i=1}^N |\psi_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i \mathbf{x}_i^\top}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} + o_P(1). \quad (\text{A.8})$$

Combining the above two equations, we have that $n^{-1} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \dot{\boldsymbol{s}}/\sqrt{n}) \mathbf{x}_i^* (\mathbf{x}_i^*)^\top$ converges in probability to $\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}$, a positive definite matrix. In addition, from Lemma 2 and Lemma 3, conditional on \mathcal{D}_N , and $\hat{\boldsymbol{\beta}}_1$, $\dot{\lambda}_r^*(\boldsymbol{\beta}_t)/\sqrt{n} = O_{P|\mathcal{D}_N, \hat{\boldsymbol{\beta}}_1}(1)$. Thus, from the Basic Corollary in page 2 of Hjort and Pollard (2011), the maximizer of $\gamma(\mathbf{s})$, $\sqrt{n}(\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_t)$, satisfies

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_r - \boldsymbol{\beta}_t) = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \frac{1}{\sqrt{n}} \dot{\lambda}_r^*(\boldsymbol{\beta}_t) + O_{P|\mathcal{D}_N, \hat{\boldsymbol{\beta}}_1}(1) \quad (\text{A.9})$$

given \mathcal{D}_N and $\hat{\boldsymbol{\beta}}_1$. Thus,

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_r - \check{\boldsymbol{\beta}}_w) = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \left\{ \frac{1}{\sqrt{n}} \dot{\lambda}_r^*(\boldsymbol{\beta}_t) - \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} \sqrt{n}(\check{\boldsymbol{\beta}}_w - \boldsymbol{\beta}_t) \right\} + O_{P|\mathcal{D}_N, \hat{\boldsymbol{\beta}}_1}(1). \quad (\text{A.10})$$

From Lemma 2,

$$\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} \sqrt{n}(\hat{\boldsymbol{\beta}}_{wf} - \boldsymbol{\beta}_t) = \frac{\sqrt{n} \sum_{i=1}^N \eta_i}{2N \mathbb{E}\{\phi(\boldsymbol{\beta}_t) h(\mathbf{x})\}} = \frac{\sqrt{n} \sum_{i=1}^N \eta_i}{N \Psi_N(\hat{\boldsymbol{\beta}}_1)} + o_P(1), \quad (\text{A.11})$$

Combining equations (A.10) and (A.11), Lemma 3, and Slutsky's theorem, Theorem 1 follows. \square

A.1.1 Proof of Proposition 1

Proof of Proposition 1. To prove that $\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \leq \mathbf{V}^{\text{OS}} = \mathbf{M}^{-1} \mathbf{V}_c^{\text{OS}} \mathbf{M}^{-1}$, we just need to show that

$$\boldsymbol{\Sigma}_{\hat{\boldsymbol{\beta}}_f}^{-1} \geq \mathbf{M}(\mathbf{V}_c^{\text{OS}})^{-1} \mathbf{M}. \quad (\text{A.12})$$

From the strong law of large numbers,

$$\mathbf{M} = \frac{1}{N} \sum_{i=1}^N \phi_i(\boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top + o(1), \quad (\text{A.13})$$

$$\mathbf{V}_c^{\text{OS}} = 4\Phi(\boldsymbol{\beta}_t) \frac{1}{N} \sum_{i=1}^N \frac{\phi_i(\boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top}{h(\mathbf{x}_i)} + o(1), \quad (\text{A.14})$$

$$\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} = 4\Phi(\boldsymbol{\beta}_t) \left\{ \frac{1}{N} \sum_{i=1}^N \phi_i(\boldsymbol{\beta}_t) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \right\}^{-1} + o(1), \quad (\text{A.15})$$

almost surely. Thus, we only need to verify that

$$\sum_{i=1}^N \phi_i(\boldsymbol{\beta}_t) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \geq \left\{ \sum_{i=1}^N \phi_i(\boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top \right\} \left\{ \sum_{i=1}^N \frac{\phi_i(\boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top}{h(\mathbf{x}_i)} \right\}^{-1} \left\{ \sum_{i=1}^N \phi_i(\boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top \right\}. \quad (\text{A.16})$$

Denote $\mathbf{Z} = \{\sqrt{\phi_1(\boldsymbol{\beta}_t)}\mathbf{x}_1, \dots, \sqrt{\phi_N(\boldsymbol{\beta}_t)}\mathbf{x}_N\}^\top$, and $\mathbf{H} = \text{diag}\{h(\mathbf{x}_1), \dots, h(\mathbf{x}_N)\}$. The inequality in (A.16) can be written as

$$\mathbf{Z}^\top \mathbf{H} \mathbf{Z} \geq \mathbf{Z}^\top \mathbf{Z} (\mathbf{Z}^\top \mathbf{H}^{-1} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{Z}, \quad (\text{A.17})$$

which is true if

$$\mathbf{H} \geq \mathbf{Z} (\mathbf{Z}^\top \mathbf{H}^{-1} \mathbf{Z})^{-1} \mathbf{Z}^\top \quad (\text{A.18})$$

Note that $(\mathbf{H}^{-1/2} \mathbf{Z}) \{(\mathbf{H}^{-1/2} \mathbf{Z})(\mathbf{H}^{-1/2} \mathbf{Z})^\top\}^{-1} (\mathbf{H}^{-1/2} \mathbf{Z})^\top$ is the projection matrix of $\mathbf{H}^{-1/2} \mathbf{Z}$, so it is, under the Loewner ordering, smaller than or equal to the identity matrix \mathbf{I}_N , namely,

$$\mathbf{I}_N \geq \mathbf{H}^{-1/2} \mathbf{Z} (\mathbf{Z}^\top \mathbf{H}^{-1} \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{H}^{-1/2}, \quad (\text{A.19})$$

which implies (A.18). If $h(\mathbf{x}) = 1$, the equality can be verified directly. \square

A.1.2 Proof of Lemma 1

Proof of Lemma 1. Let B be a bound for g_1 i.e., $|g_1| \leq B$. For any $\epsilon > 0$, by Markov's inequality,

$$\begin{aligned} \mathbb{P}\left\{\left|\frac{1}{n} \sum_{i=1}^n g_1(\mathbf{v}_i) g_2(\mathbf{v}_i)\right| > \epsilon\right\} &\leq \frac{\mathbb{E}|g_1(\mathbf{v}) g_2(\mathbf{v})|}{\epsilon} \\ &= \frac{\mathbb{E}[|g_1(\mathbf{v})| |g_2(\mathbf{v})| I\{|g_2(\mathbf{v})| \leq K\}]}{\epsilon} + \frac{\mathbb{E}[|g_1(\mathbf{v})| |g_2(\mathbf{v})| I\{|g_2(\mathbf{v})| > K\}]}{\epsilon} \\ &\leq \frac{K}{\epsilon} \mathbb{E}|g_1(\mathbf{v})| + \frac{B}{\epsilon} \mathbb{E}\{|g_2(\mathbf{v})| I\{|g_2(\mathbf{v})| > K\}\}. \end{aligned}$$

For any $\zeta > 0$, we can choose a K large enough such that $\mathbb{E}\{|g_2(\mathbf{v})| I\{|g_2(\mathbf{v})| \leq K\}\} < \zeta \epsilon / (2B)$, since $\mathbb{E}|g_2(\mathbf{v})| < \infty$. The facts that $g_1(\mathbf{v}_i) \leq B$ and $g_1(\mathbf{v}_i) = o_P(1)$ imply that $\mathbb{E}|g_1(\mathbf{v})| = o(1)$. Thus, there is a n_ζ such that $\mathbb{E}|g_1(\mathbf{v})| < \zeta \epsilon / (2K)$ when $n > n_\zeta$. Therefore, for any $\zeta > 0$, $\mathbb{P}\{|n^{-1} \sum_{i=1}^n g_1(\mathbf{v}_i) g_2(\mathbf{v}_i)| > \epsilon\} < \zeta$ for sufficiently large n . This finishes the proof. \square

A.1.3 Proof of Lemma 2

Proof of Lemma 2. Since $\check{\boldsymbol{\beta}}_w$ is the maximizer of

$$\lambda_{wf}(\boldsymbol{\beta}) = \sum_{i=1}^N |y_i - p_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) [y_i \mathbf{x}_i^\top (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1) - \log\{1 + e^{\mathbf{x}_i^\top (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)}\}],$$

$\sqrt{N}(\check{\boldsymbol{\beta}}_w - \boldsymbol{\beta}_t)$ is the maximizer of $\gamma_{wf}(\mathbf{s}) = \lambda_{wf}(\boldsymbol{\beta}_t + \mathbf{s}/\sqrt{N}) - \lambda_{wf}(\boldsymbol{\beta}_t)$. By Taylor's expansion,

$$\gamma_{wf}(\mathbf{s}) = \frac{1}{\sqrt{N}} \mathbf{s}^\top \dot{\lambda}_{wf}(\boldsymbol{\beta}_t) - \frac{1}{2N} \sum_{i=1}^N |y_i - p_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}/\sqrt{N}) (\mathbf{s}^\top \mathbf{x}_i)^2 \quad (\text{A.20})$$

where

$$\dot{\lambda}_{wf}(\boldsymbol{\beta}_t) = \sum_{i=1}^N \eta_i = \sum_{i=1}^N |\psi_i(\hat{\boldsymbol{\beta}}_1)| \psi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}_i) \mathbf{x}_i, \quad (\text{A.21})$$

and \acute{s} lies between $\mathbf{0}$ and \mathbf{s} .

Denote $\psi(\boldsymbol{\beta}) = y - p(\boldsymbol{\beta})$ and $\eta = |\psi(\hat{\boldsymbol{\beta}}_1)| \psi(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}) \mathbf{x}$. Note that $\mathbb{E}(\eta|\hat{\boldsymbol{\beta}}_1) = \mathbf{0}$ because

$$\begin{aligned} \mathbb{E}(\eta|\hat{\boldsymbol{\beta}}_1, \mathbf{x}) &= \mathbb{E} \left[|\psi(\hat{\boldsymbol{\beta}}_1)| \{y - p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\} h(\mathbf{x}) \mathbf{x} \middle| \hat{\boldsymbol{\beta}}_1, \mathbf{x} \right] \\ &= -p(\mathbf{x}, \hat{\boldsymbol{\beta}}_1) p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \{1 - p(\mathbf{x}, \boldsymbol{\beta}_t)\} h(\mathbf{x}) \mathbf{x} \\ &\quad + \{1 - p(\mathbf{x}, \hat{\boldsymbol{\beta}}_1)\} \{1 - p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\} p(\mathbf{x}, \boldsymbol{\beta}_t) h(\mathbf{x}) \mathbf{x} \\ &= -\frac{e^{\mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}}{1 + e^{\mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}} \frac{e^{\mathbf{x}^\top \boldsymbol{\beta}_t - \mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}}{1 + e^{\mathbf{x}^\top \boldsymbol{\beta}_t - \mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}} \frac{1}{1 + e^{\mathbf{x}^\top \boldsymbol{\beta}_t}} h(\mathbf{x}) \mathbf{x} \\ &\quad + \frac{1}{1 + e^{\mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}} \frac{1}{1 + e^{\mathbf{x}^\top \boldsymbol{\beta}_t - \mathbf{x}^\top \hat{\boldsymbol{\beta}}_1}} \frac{e^{\mathbf{x}^\top \boldsymbol{\beta}_t}}{1 + e^{\mathbf{x}^\top \boldsymbol{\beta}_t}} h(\mathbf{x}) \mathbf{x} = 0. \end{aligned} \quad (\text{A.22})$$

This also gives that

$$\mathbb{V}(\eta|\hat{\boldsymbol{\beta}}_1) = \mathbb{E}\{\mathbb{V}(\eta|\mathbf{x}, \hat{\boldsymbol{\beta}}_1)|\hat{\boldsymbol{\beta}}_1\} + \mathbb{V}\{\mathbb{E}(\eta|\mathbf{x}, \hat{\boldsymbol{\beta}}_1)|\hat{\boldsymbol{\beta}}_1\} = \mathbb{E}\{\mathbb{V}(\eta|\mathbf{x}, \hat{\boldsymbol{\beta}}_1)|\hat{\boldsymbol{\beta}}_1\}.$$

Now, since

$$\begin{aligned} \mathbb{V}(\eta|\mathbf{x}, \hat{\boldsymbol{\beta}}_1) &= \mathbb{E} \left[|y - p(\hat{\boldsymbol{\beta}}_1)|^2 \{y - p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\}^2 h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top \middle| \hat{\boldsymbol{\beta}}_1, \mathbf{x} \right] \\ &= p(\mathbf{x}, \boldsymbol{\beta}_t) \{1 - p(\hat{\boldsymbol{\beta}}_1)\}^2 \{1 - p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\}^2 h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top \\ &\quad + \{1 - p(\mathbf{x}, \boldsymbol{\beta}_t)\} \{p(\hat{\boldsymbol{\beta}}_1)\}^2 \{p(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\}^2 h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top \\ &= \phi_i(\hat{\boldsymbol{\beta}}_1) \phi_i(\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top, \end{aligned}$$

we have

$$\mathbb{V}(\eta|\hat{\boldsymbol{\beta}}_1) = \mathbb{E} \left\{ \phi(\hat{\boldsymbol{\beta}}_1) \phi(\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top \middle| \hat{\boldsymbol{\beta}}_1 \right\}. \quad (\text{A.23})$$

Let $\| \cdot \|$ denote the Frobenius norm if applied on a martix, i.e., for a matrix A , $\|A\|^2 = \text{tr}(AA^\top)$, and denote $\mathbb{V}(\eta|\boldsymbol{\beta}_t) = \mathbb{E}\{\phi(\boldsymbol{\beta}_t) \phi(\boldsymbol{\beta}_t - \boldsymbol{\beta}_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\} = 0.25 \mathbb{E}\{\phi(\boldsymbol{\beta}_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\}$. Notice that $|\phi(\hat{\boldsymbol{\beta}}_1) \phi(\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_t) - 0.25 \phi_i(\boldsymbol{\beta}_t)| h^2(\mathbf{x}) \|\mathbf{x}\|^2$ converges to 0 in probability and it is bounded by $h^2(\mathbf{x}) \|\mathbf{x}\|^2$, an integrable random variable under Assumption 2. Thus,

$$\mathbb{E} \|\mathbb{V}(\eta|\hat{\boldsymbol{\beta}}_1) - \mathbb{V}(\eta|\boldsymbol{\beta}_t)\| \leq \mathbb{E} \{ |\phi(\hat{\boldsymbol{\beta}}_1) \phi(\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_t) - 0.25 \phi_i(\boldsymbol{\beta}_t)| h^2(\mathbf{x}) \|\mathbf{x}\|^2 \} = o(1). \quad (\text{A.24})$$

This implies that

$$\mathbb{V}(\eta|\hat{\boldsymbol{\beta}}_1) = \mathbb{V}(\eta|\boldsymbol{\beta}_t) + o_P(1) = 0.25 \mathbb{E}\{\phi(\boldsymbol{\beta}_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\} + o_P(1). \quad (\text{A.25})$$

Conditional on $\hat{\beta}_1$, η_1, \dots, η_N are i.i.d. with mean $\mathbf{0}$ and variance $\mathbb{V}(\eta|\hat{\beta}_1)$. Since for any $\epsilon > 0$,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left\{ \|\eta_i\|^2 I(\|\eta_i\| > \sqrt{N}\epsilon) \middle| \hat{\beta}_1 \right\} &\leq \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left\{ \|h(\mathbf{x}_i)\mathbf{x}_i\|^2 I(\|h(\mathbf{x}_i)\mathbf{x}_i\| > \sqrt{N}\epsilon) \middle| \hat{\beta}_1 \right\} \\ &= \mathbb{E} \left\{ \|h(\mathbf{x})\mathbf{x}\|^2 I(\|h(\mathbf{x})\mathbf{x}\| > \sqrt{N}\epsilon) \right\} \rightarrow 0, \end{aligned}$$

the Lindeberg-Feller central limit theorem (Section *2.8 of van der Vaart, 1998) applies conditional on $\hat{\beta}_1$. Thus, we have, conditional on $\hat{\beta}_1$,

$$\frac{\dot{\lambda}_{wf}(\beta_t)}{\sqrt{N}} \longrightarrow \mathbb{N} \left[\mathbf{0}, \frac{\mathbb{E} \{ \phi(\beta_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^T \}}{4} \right], \quad (\text{A.26})$$

From Lemma 1, conditional on $\hat{\beta}_1$,

$$\begin{aligned} &\frac{1}{N} \sum_{i=1}^N |y_i - p_i(\hat{\beta}_1)| h(\mathbf{x}_i) \phi_i(\beta_t - \hat{\beta}_1 + \dot{s}/\sqrt{N}) \mathbf{x}_i \mathbf{x}_i^T \\ &= \frac{1}{4} \mathbb{E} \{ |\psi(\beta_t)| h(\mathbf{x}) \mathbf{x} \mathbf{x}^T \} + o_{P|\hat{\beta}_1}(1) = \frac{1}{2} \mathbb{E} \{ \phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^T \} + o_P(1). \end{aligned}$$

Thus, from the Basic Corollary in page 2 of Hjort and Pollard (2011), the maximizer of $\gamma_{wf}(\mathbf{s})$, $\sqrt{N}(\hat{\beta}_w - \beta_t)$, satisfies

$$\sqrt{N}(\hat{\beta}_w - \beta_t) = 2[\mathbb{E} \{ \phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^T \}]^{-1} \frac{1}{\sqrt{N}} \dot{\lambda}_w(\beta_t) + o_P(1). \quad (\text{A.27})$$

Note that

$$[\mathbb{E} \{ \phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^T \}]^{-1} = \frac{\Sigma_{\beta_t}}{4\Phi(\beta_t)}. \quad (\text{A.28})$$

Combining equations (A.21), (A.27), and (A.28), we have

$$\sqrt{N}(\hat{\beta}_w - \beta_t) = \frac{\Sigma_{\beta_t}}{2\Phi(\beta_t)} \frac{1}{\sqrt{N}} \dot{\lambda}_w(\beta_t) + o_P(1).$$

An application of Slutsky's theorem yields the result for the asymptotic normality. \square

A.1.4 Proof of Lemma 3

Proof of Lemma 3. Note that given \mathcal{D}_N and $\hat{\beta}_1$, $\{y_i^* - p_i^*(\beta_t - \hat{\beta}_1)\} \mathbf{x}_i^*$ are i.i.d. random vectors. We now exam their mean and variance, and check the Lindeberg-Feller condition (Section *2.8 of van der Vaart, 1998) under the conditional distribution given \mathcal{D}_N and $\hat{\beta}_1$. For the expectation, we have,

$$\mathbb{E} [\{y^* - p^*(\beta - \hat{\beta}_1)\} \mathbf{x}^* | \mathcal{D}_N, \hat{\beta}_1] = \sum_{i=1}^N \pi_i(\hat{\beta}_1) \psi_i(\beta_t - \hat{\beta}_1) \mathbf{x}_i = \frac{\sum_{i=1}^N \eta_i}{N \Psi_N(\hat{\beta}_1)}. \quad (\text{A.29})$$

From Lemma 2 and its proof, $\sum_{i=1}^N \eta_i = O_P(\sqrt{N})$ conditional on $\hat{\boldsymbol{\beta}}_1$ in probability, i.e., for any $\epsilon > 0$, there exists a K such that $\mathbb{P}\{\mathbb{P}(\sum_{i=1}^N \eta_i/\sqrt{N} > K|\hat{\boldsymbol{\beta}}_1) < \epsilon\} \rightarrow 1$ as $n_1, N \rightarrow \infty$. From Xiong and Li (2008), we know that $\sum_{i=1}^N \eta_i = O_P(\sqrt{N})$ unconditionally. Thus, for the expectation, we have

$$\Delta = \mathbb{E}[\{y^* - p^*(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)\} \mathbf{x}^* | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1] = O_P(1/\sqrt{N}).$$

For the variance,

$$\begin{aligned} & \mathbb{V}[\{y^* - p^*(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)\} \mathbf{x}^* | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1] \\ &= \sum_{i=1}^N \pi_i \{y_i - p_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\}^2 \mathbf{x}_i \mathbf{x}_i^\top - \Delta^2 \\ &= \frac{1}{N} \sum_{i=1}^N \frac{|\psi_i(\hat{\boldsymbol{\beta}}_1)| \psi_i^2(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} - O_P(1/N) \\ &= \frac{1}{N} \sum_{i=1}^N \frac{|\psi_i(\boldsymbol{\beta}_t)| (y_i - 0.5)^2 h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top}{\Psi_N(\boldsymbol{\beta}_t)} + o_P(1) \\ &= \frac{1}{4} \frac{\mathbb{E}\{|\psi(\boldsymbol{\beta}_t)| h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\}}{\Psi(\boldsymbol{\beta}_t)} + o_P(1) = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} + o_P(1) \end{aligned}$$

where the third equality is from Lemma 1 and the fact that $\mathbb{E}\{h(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$, and the fourth equality is from the law of large numbers.

Now we check the Lindeberg-Feller condition (Section *2.8 of van der Vaart, 1998) under the condition distribution. Denote $\dot{\lambda}_{ri}^* = \{y_i^* - p_i^*(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)\} \mathbf{x}_i^*$.

$$\mathbb{E} \frac{1}{n} \sum_{i=1}^n \{ \|\dot{\lambda}_{ri}^*\|^2 I(\|\dot{\lambda}_{ri}^*\| > \sqrt{n}\epsilon) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1 \} \quad (\text{A.30})$$

$$\leq \mathbb{E} \{ \|\mathbf{x}^*\|^2 I(\|\mathbf{x}\| > \sqrt{n}\epsilon) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1 \} \quad (\text{A.31})$$

$$= \sum_{i=1}^N \pi(\hat{\boldsymbol{\beta}}_1) \{ \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\| > \sqrt{n}\epsilon) \} \quad (\text{A.32})$$

$$\leq \frac{1}{N} \sum_{i=1}^N \frac{\{h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\| > \sqrt{n}\epsilon)\}}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} \quad (\text{A.33})$$

$$\leq \frac{1}{N} \sum_{i=1}^N \frac{\{h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\| > \sqrt{n}\epsilon)\}}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} = o_P(1), \quad (\text{A.34})$$

by Lemma 1 and the fact that $\mathbb{E}\{h(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$. Thus, applying the Lindeberg-Feller central limit theorem (Section *2.8 of van der Vaart, 1998) finishes the proof. \square

A.1.5 Proof of Lemma 4

Proof of Lemma 4. We begin with the following partition,

$$\frac{1}{n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i^*\|^2 \quad (\text{A.35})$$

$$\begin{aligned} &= \frac{1}{n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i^*\|^2 I(\|\mathbf{x}_i^*\|^2 \leq n) + \frac{1}{n} \sum_{i=1}^n \phi_i^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i^*\|^2 I(\|\mathbf{x}_i^*\|^2 > n) \\ &\equiv \Delta_1 + \Delta_2. \end{aligned} \quad (\text{A.36})$$

The second term Δ_2 is $o_P(1)$ because it is non-negative and

$$\mathbb{E}(\Delta_2 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) = \sum_{i=1}^N \pi_i(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\|^2 > n) \quad (\text{A.37})$$

$$\leq \frac{\sum_{i=1}^N |\psi_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\|^2 > n)}{\sum_{i=1}^N |\psi_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i)} \quad (\text{A.38})$$

$$\leq \frac{\frac{1}{N} \sum_{i=1}^N h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\|\mathbf{x}_i\|^2 > n)}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} = o_P(1) \quad (\text{A.39})$$

as $n, N \rightarrow \infty$, where the last step is from Lemma 1.

Similarly, we can show that

$$\mathbb{E}(\Delta_1 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) - \sum_{i=1}^N \pi(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \|\mathbf{x}_i\|^2 = o_P(1). \quad (\text{A.40})$$

Thus, we only need to show that $\Delta_1 - \mathbb{E}(\Delta_1 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) = o_P(1)$. For this, we show that the conditional variance of Δ_1 goes to 0 in probability. Notice that

$$\begin{aligned} &\mathbb{V}(\Delta_1 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) \\ &= \frac{1}{n} \mathbb{V}\{\phi^*(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \|\mathbf{x}^*\|^2 I(\|\mathbf{x}^*\|^2 \leq n) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1\} \\ &\leq \frac{\|\mathbf{s}\|^4}{16n} \mathbb{E}\{\|\mathbf{x}^*\|^4 I(\|\mathbf{x}^*\|^2 \leq n) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1\} \\ &= \frac{\|\mathbf{s}\|^4}{16n} \sum_{i=1}^n \mathbb{E}\{\|\mathbf{x}^*\|^4 I(i-1 < \|\mathbf{x}^*\|^2 \leq i) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1\} \\ &\leq \frac{\|\mathbf{s}\|^4}{16n} \sum_{i=1}^n i^2 \mathbb{E}\{I(i-1 < \|\mathbf{x}^*\|^2 \leq i) | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1\} \\ &\leq \frac{\|\mathbf{s}\|^4}{16n} \sum_{i=1}^n i^2 \{\mathbb{P}(\|\mathbf{x}^*\|^2 > i-1 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) - \mathbb{P}(\|\mathbf{x}^*\|^2 > i | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1)\} \end{aligned}$$

$$\begin{aligned}
&= \frac{\|\mathbf{s}\|^4}{16n} \left\{ \mathbb{P}(\|\mathbf{x}^*\|^2 > 0 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) - n^2 \mathbb{P}(\|\mathbf{x}^*\|^2 > n | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) + \sum_{i=1}^{n-1} (2i+1) \mathbb{P}(\|\mathbf{x}^*\|^2 > i | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) \right\} \\
&\leq \frac{\|\mathbf{s}\|^4}{16n} \left\{ 1 + \sum_{i=1}^{n-1} 3i \mathbb{P}(\|\mathbf{x}^*\|^2 > i | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) \right\}
\end{aligned}$$

This is $o_P(1)$ because

$$\frac{1}{n} \sum_{i=1}^n i \mathbb{P}(\|\mathbf{x}^*\|^2 > i | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1) = \frac{1}{n} \sum_{i=1}^n i \sum_{j=1}^n \pi_j(\hat{\boldsymbol{\beta}}_1) I(\|\mathbf{x}_j\|^2 > i) \quad (\text{A.41})$$

$$= \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^n i \pi_j(\hat{\boldsymbol{\beta}}_1) I(\|\mathbf{x}_j\|^2 > i) \quad (\text{A.42})$$

$$= \frac{\frac{1}{N} \sum_{j=1}^n \frac{1}{n} \sum_{i=1}^n i |\psi_j(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_j) I(\|\mathbf{x}_j\|^2 > i)}{\Psi_N(\hat{\boldsymbol{\beta}}_1)} \quad (\text{A.43})$$

$$\leq \frac{\frac{1}{N} \sum_{j=1}^n \frac{1}{n} \sum_{i=1}^n i h(\mathbf{x}_j) I(\|\mathbf{x}_j\|^2 > i)}{\Psi_N(\hat{\boldsymbol{\beta}}_1)}, \quad (\text{A.44})$$

and the numerator is non-negative and has an expectation

$$\frac{1}{N} \sum_{j=1}^n \frac{1}{n} \sum_{i=1}^n i \mathbb{E}\{h(\mathbf{x}) I(\|\mathbf{x}\|^2 > i)\} \quad (\text{A.45})$$

which is $o(1)$ since $i \mathbb{E}\{h(\mathbf{x}) I(\|\mathbf{x}\|^2 > i)\} = o(1)$ as $i \rightarrow \infty$. \square

A.2 Proofs for Poisson subsampling

In this section we prove the results in Section 4 about Poisson subsampling.

Define $\delta_i^{\hat{\boldsymbol{\beta}}_1} = I\{u_i \leq n\pi_i^p(\hat{\boldsymbol{\beta}}_1)\}$, and use notation λ_p to denote the log-likelihood shifted by $\hat{\boldsymbol{\beta}}_1$, i.e., $\lambda_p(\boldsymbol{\beta}) = \ell_p^*(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)$. Using these notation, the estimator $\hat{\boldsymbol{\beta}}_p$ is the maximizer of

$$\lambda_p(\boldsymbol{\beta}) = \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} [(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^T \mathbf{x}_i y_i - \log\{1 + e^{(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^T \mathbf{x}_i}\}], \quad (\text{A.46})$$

Denote the first and second derivatives of $\lambda_p(\boldsymbol{\beta})$ as $\dot{\lambda}_p(\boldsymbol{\beta}) = \partial \lambda_p(\boldsymbol{\beta}) / \partial \boldsymbol{\beta}$ and $\ddot{\lambda}_p(\boldsymbol{\beta}) = \partial^2 \lambda_p(\boldsymbol{\beta}) / (\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}^T)$. Two lemmas similar to Lemmas 3 and 4 are derived below which will be used in the proof of Theorem 2.

Lemma 5. *Let*

$$\dot{\lambda}_p(\boldsymbol{\beta}_t) = \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \{y_i - p_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\} \mathbf{x}_i. \quad (\text{A.47})$$

Under Assumptions 1 and 2, conditional on \mathcal{D}_N , $\hat{\boldsymbol{\beta}}_1$, and $\hat{\Psi}_1$, if $n = o(N)$, then

$$\frac{\dot{\lambda}_p(\boldsymbol{\beta}_t)}{\sqrt{n}} - \frac{\sqrt{n} \sum_{i=1}^N \eta_i}{N \Psi_N(\hat{\boldsymbol{\beta}}_1)} \longrightarrow \mathbb{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}), \quad (\text{A.48})$$

in distribution; if $n/N \rightarrow \rho \in (0, 1)$, then

$$\frac{\dot{\lambda}_p(\boldsymbol{\beta}_t)}{\sqrt{n}} - \frac{\sqrt{n} \sum_{i=1}^N \eta_i}{N \Psi_N(\hat{\boldsymbol{\beta}}_1)} \longrightarrow \mathbb{N}(\mathbf{0}, \boldsymbol{\Lambda}_\rho), \quad (\text{A.49})$$

in distribution.

Lemma 6. Under Assumptions 1 and 2, as $n, N \rightarrow \infty$, for any $\mathbf{s}_n \rightarrow \mathbf{0}$,

$$\frac{1}{n} \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) \|\mathbf{x}_i\|^2 - \sum_{i=1}^N \pi_i^p(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \|\mathbf{x}_i\|^2 = o_P(1).$$

Proof of Theorem 2. The estimator $\hat{\boldsymbol{\beta}}_p$ is the maximizer of (A.46), so $\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t)$ is the maximizer of $\gamma_p(\mathbf{s}) = \lambda_p(\boldsymbol{\beta}_t + \mathbf{s}/\sqrt{n}) - \lambda_p(\boldsymbol{\beta}_t)$. By Taylor's expansion,

$$\gamma_p(\mathbf{s}) = \frac{1}{\sqrt{n}} \mathbf{s}^\top \dot{\lambda}_p(\boldsymbol{\beta}_t) + \frac{1}{2n} \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}/\sqrt{n}) (\mathbf{s}^\top \mathbf{x}_i)^2 \quad (\text{A.50})$$

where $\phi_i(\boldsymbol{\beta}) = p_i(\boldsymbol{\beta})\{1 - p_i(\boldsymbol{\beta})\}$, and \mathbf{s} lies between $\mathbf{0}$ and \mathbf{s} .

From Lemmas 5 and 6, conditional on \mathcal{D}_N , and $\hat{\boldsymbol{\beta}}_1$,

$$\frac{1}{n} \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}/\sqrt{n}) \mathbf{x}_i \mathbf{x}_i^\top = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} + o_P(1). \quad (\text{A.51})$$

In addition, from Lemma 5, conditional on \mathcal{D}_N , $\hat{\boldsymbol{\beta}}_1$, and $\hat{\Psi}_1$, $\dot{\lambda}_p(\boldsymbol{\beta}_t)/\sqrt{n}$ converges in distribution to a normal limit. Thus, from the Basic Corollary in page 2 of Hjort and Pollard (2011), the maximizer of $\gamma_p(\mathbf{s})$, $\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t)$, satisfies

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t) = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \frac{1}{\sqrt{n}} \dot{\lambda}_p(\boldsymbol{\beta}_t) + o_P(1) \quad (\text{A.52})$$

given \mathcal{D}_N , $\hat{\boldsymbol{\beta}}_1$, and $\hat{\Psi}_1$. Combining this with Lemma 5 and Slutsky's theorem, Theorem 2 follows. \square

A.2.1 Proof of Proposition 2

Proof of Proposition 2. To prove that $\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \boldsymbol{\Lambda}_\rho \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} < \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}$, we just need to show that $\boldsymbol{\Lambda}_\rho < \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}$. This is true because

$$\begin{aligned} \boldsymbol{\Lambda}_\rho &= \frac{\mathbb{E}[\phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \{\Phi(\boldsymbol{\beta}_t) - \rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x})\}_+ \mathbf{x} \mathbf{x}^\top]}{4\Phi^2(\boldsymbol{\beta}_t)} \\ &\leq \frac{\mathbb{E}\{\phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \Phi(\boldsymbol{\beta}_t) \mathbf{x} \mathbf{x}^\top\}}{4\Phi^2(\boldsymbol{\beta}_t)} = \frac{\mathbb{E}\{\phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\}}{4\Phi(\boldsymbol{\beta}_t)} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}. \end{aligned}$$

\square

A.2.2 Proof of Lemma 5

Proof of Lemma 5. Note that, $\delta_i^{\hat{\beta}_1} = I\{u_i \leq n\pi_i^p(\hat{\beta}_1)\}$, where u_i are i.i.d. with the standard uniform distribution. Thus, given \mathcal{D}_N , $\hat{\beta}_1$, and $\hat{\Psi}_1$, $\dot{\lambda}_p(\beta_t)$ is a sum of N independent random vectors. We now exam the mean and variance of $\dot{\lambda}_p(\beta_t)$. Recall that $\eta_i = |\psi_i(\hat{\beta}_1)|\psi_i(\beta_t - \hat{\beta}_1)h(\mathbf{x}_i)\mathbf{x}_i$, and $\psi_i(\beta) = y_i - p_i(\beta)$. For the mean, we have,

$$\frac{1}{\sqrt{n}}\mathbb{E}\{\dot{\lambda}_p(\beta_t)|\mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1\} \quad (\text{A.53})$$

$$= \frac{1}{\sqrt{n}}\sum_{i=1}^N \{n\pi_i^p(\hat{\beta}_1) \wedge 1\} \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i(\beta_t - \hat{\beta}_1)\mathbf{x}_i \quad (\text{A.54})$$

$$= \frac{1}{\sqrt{n}}\sum_{i=1}^N n\pi_i^p(\hat{\beta}_1)\psi_i(\beta_t - \hat{\beta}_1)\mathbf{x}_i = \frac{\sqrt{n}}{\sqrt{N}}\frac{\sum_{i=1}^N \eta_i}{\hat{\Psi}_1\sqrt{N}} = O_P(\sqrt{n/N}), \quad (\text{A.55})$$

where the last equality is from Lemma 2.

For the variance,

$$\begin{aligned} & \frac{1}{n}\mathbb{V}\{\dot{\lambda}_p(\beta_t)|\mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1\} \\ &= \frac{1}{n}\sum_{i=1}^N [\{n\pi_i^p(\hat{\beta}_1) \wedge 1\} - \{n\pi_i^p(\hat{\beta}_1) \wedge 1\}^2] \{n\pi_i^p(\hat{\beta}_1) \vee 1\}^2 \psi_i^2(\beta_t - \hat{\beta}_1)\mathbf{x}_i\mathbf{x}_i^T \\ &= \sum_{i=1}^N \pi_i^p(\hat{\beta}_1) \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i^2(\beta_t - \hat{\beta}_1)\mathbf{x}_i\mathbf{x}_i^T - n \sum_{i=1}^N \{\pi_i^p(\hat{\beta}_1)\}^2 \psi_i^2(\beta_t - \hat{\beta}_1)\mathbf{x}_i\mathbf{x}_i^T \\ &= \frac{1}{N} \sum_{i=1}^N |\psi_i(\hat{\beta}_1)| \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i^2(\beta_t - \hat{\beta}_1) h(\mathbf{x}_i)\mathbf{x}_i\mathbf{x}_i^T \\ &= \frac{\hat{\Psi}_1}{N} \frac{n \frac{1}{N} \sum_{i=1}^N \psi_i^2(\hat{\beta}_1) \psi_i^2(\beta_t - \hat{\beta}_1) h^2(\mathbf{x}_i)\mathbf{x}_i\mathbf{x}_i^T}{\hat{\Psi}_1^2} \\ &\equiv \Delta_3 + \Delta_4 \end{aligned} \quad (\text{A.56})$$

Note that $\mathbb{E}\{h(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$, $\mathbb{E}\{h^2(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$, and $|\psi_i(\cdot)|$ are bounded. Thus, from Lemma 1, if $n/N \rightarrow \rho$,

$$\Delta_4 \rightarrow \rho \frac{\mathbb{E}\{\psi^2(\beta_t)h^2(\mathbf{x})\mathbf{x}\mathbf{x}^T\}}{4\Psi^2(\beta_t)} = \rho \frac{\mathbb{E}\{\phi^2(\beta_t)h^2(\mathbf{x})\mathbf{x}\mathbf{x}^T\}}{4\Phi^2(\beta_t)}, \quad (\text{A.57})$$

in probability.

For the term Δ_3 in (A.56), it is equal to

$$\Delta_3 = \frac{1}{\hat{\Psi}_1^2} \frac{1}{N} \sum_{i=1}^N |\psi_i(\hat{\beta}_1)| \left\{ \frac{n|\psi_i(\hat{\beta}_1)|h(\mathbf{x}_i)}{N} \vee \hat{\Psi}_1 \right\} \psi_i^2(\beta_t - \hat{\beta}_1) h(\mathbf{x}_i)\mathbf{x}_i\mathbf{x}_i^T$$

$$\begin{aligned}
&= \frac{1}{\hat{\Psi}_1^2} \frac{n}{N^2} \sum_{i=1}^N \psi_i^2(\hat{\beta}_1) \psi_i^2(\beta_t - \hat{\beta}_1) h^2(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top I \left\{ \frac{n|\psi_i(\hat{\beta}_1)|h(\mathbf{x}_i)}{N} > \hat{\Psi}_1 \right\} \\
&\quad + \frac{1}{\hat{\Psi}_1} \frac{1}{N} \sum_{i=1}^N |\psi_i(\hat{\beta}_1)| \psi_i^2(\beta_t - \hat{\beta}_1) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top I \left\{ \frac{n|\psi_i(\hat{\beta}_1)|h(\mathbf{x}_i)}{N} \leq \hat{\Psi}_1 \right\}.
\end{aligned}$$

Since $\mathbb{E}\{h(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$, $\mathbb{E}\{h^2(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$, and $|\psi_i(\cdot)|$ are bounded, from Lemma 1, if $n/N \rightarrow \rho$, as $n \rightarrow \infty$ and $N \rightarrow \infty$, in probability,

$$\begin{aligned}
\Delta_3 &\rightarrow \frac{\rho \mathbb{E}[\psi^2(\beta_t) h^2(\mathbf{x}) \mathbf{x} \mathbf{x}^\top I \{\rho |\psi(\beta_t)| h(\mathbf{x}) \geq \Psi(\beta_t)\}]}{4\Psi^2(\beta_t)} \\
&\quad + \frac{\mathbb{E}[|\psi(\beta_t)| h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top I \{\rho |\psi(\beta_t)| h(\mathbf{x}) \leq \Psi(\beta_t)\}]}{4\Psi(\beta_t)} \\
&= \frac{\mathbb{E}(\phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top [\{\rho \phi(\beta_t) h(\mathbf{x})\} \vee \Phi(\beta_t)])}{4\Phi^2(\beta_t)} \tag{A.58}
\end{aligned}$$

From, (A.56), (A.57), and (A.58), if $n/N \rightarrow \rho$,

$$\frac{1}{n} \mathbb{V}\{\dot{\lambda}_p(\beta_t) | \mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1\} = \frac{\mathbb{E}[\phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top \{\Phi(\beta_t) - \rho \phi(\beta_t) h(\mathbf{x})\}_+]}{4\Phi^2(\beta_t)} + o_P(1). \tag{A.59}$$

Specifically, when $\rho = 0$,

$$\frac{1}{n} \mathbb{V}\{\dot{\lambda}_p(\beta_t) | \mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1\} = \Sigma_{\beta_t} + o_P(1). \tag{A.60}$$

Now we check the Lindeberg-Feller condition (Section *2.8 of van der Vaart, 1998) under the condition distribution. Denote $\dot{\lambda}_{pi} = \delta_i^{\hat{\beta}_1} \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i(\beta_t - \hat{\beta}_1) \mathbf{x}_i$. For any $\epsilon > 0$

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^N \mathbb{E}\left\{ \|\dot{\lambda}_{pi}\|^2 I(\|\dot{\lambda}_{pi}\| > \sqrt{n}\epsilon) \middle| \mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1 \right\} \\
&\leq \frac{1}{n} \sum_{i=1}^N \mathbb{E}\left[\|\delta_i^{\hat{\beta}_1} \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \mathbf{x}_i\|^2 I(\|\delta_i^{\hat{\beta}_1} \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \mathbf{x}_i\| > \sqrt{n}\epsilon) \middle| \mathcal{D}_N, \hat{\beta}_1, \hat{\Psi}_1 \right] \\
&= \sum_{i=1}^N \pi_i^p(\hat{\beta}_1) \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \|\mathbf{x}_i\|^2 I(\{n\pi_i^p(\hat{\beta}_1) \vee 1\} \|\mathbf{x}_i\| > \sqrt{n}\epsilon) \\
&\leq \frac{|\psi_i(\hat{\beta}_1)| h(\mathbf{x}_i) \{n/N |\psi_i(\hat{\beta}_1)| h(\mathbf{x}_i) + \hat{\Psi}_1\} \|\mathbf{x}_i\|^2 I(\{n\pi_i^p(\hat{\beta}_1) + 1\} \|\mathbf{x}_i\| > \sqrt{n}\epsilon)}{\hat{\Psi}_1^2} \\
&\leq \frac{\frac{1}{N} \sum_{i=1}^N h^2(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\{h(\mathbf{x}_i)/\hat{\Psi}_1 + 1\} \|\mathbf{x}_i\| > \sqrt{n}\epsilon)}{\hat{\Psi}_1^2} \\
&\quad + \frac{\frac{1}{N} \sum_{i=1}^N h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 I(\{h(\mathbf{x}_i)/\hat{\Psi}_1 + 1\} \|\mathbf{x}_i\| > \sqrt{n}\epsilon)}{\hat{\Psi}_1} = o_P(1),
\end{aligned}$$

where the last equality is from Lemma 1. Thus, applying the Lindeberg-Feller central limit theorem (Section *2.8 of van der Vaart, 1998) finishes the proof. \square

A.2.3 Proof of Lemma 6

Proof of Lemma 6. Note that, from Lemma 1,

$$\sum_{i=1}^N \pi_i^p(\hat{\boldsymbol{\beta}}_1) \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i \mathbf{x}_i^\top = \frac{1}{\hat{\Psi}_1 N} \sum_{i=1}^N |\psi_i(\hat{\boldsymbol{\beta}}_1)| \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} + o_P(1);$$

and from the strong law of large numbers

$$\frac{1}{n} \sum_{i=1}^N \delta_i^{\boldsymbol{\beta}_t} \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_t) \mathbf{x}_i \mathbf{x}_i^\top = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} + o_P(1),$$

where $\delta_i^{\boldsymbol{\beta}_t} = I\{u_i \leq n\pi_i^p(\boldsymbol{\beta}_t)\}$. Thus, if we show that

$$\Delta_5 \equiv \frac{1}{n} \sum_{i=1}^N \left| \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) - \delta_i^{\boldsymbol{\beta}_t} \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_t) \right| \|\mathbf{x}_i\|^2 = o_P(1), \quad (\text{A.61})$$

then the result in Lemma 6 follows. Noting that Δ_5 is nonnegative, we prove (A.61) by showing that $\mathbb{E}(\Delta_5 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) = o_P(1)$. Note that given \mathcal{D}_N , $\hat{\boldsymbol{\beta}}_1$, and $\hat{\Psi}_1$, the only random terms in Δ_5 are $\delta_i^{\hat{\boldsymbol{\beta}}_1} = I\{u_i \leq n\pi_i^p(\hat{\boldsymbol{\beta}}_1)\}$ and $\delta_i^{\boldsymbol{\beta}_t} = I\{u_i \leq n\pi_i^p(\boldsymbol{\beta}_t)\}$. We have that

$$\begin{aligned} & \mathbb{E}(\Delta_5 | \mathcal{D}_N, \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) \\ & \leq \frac{1}{n} \sum_{i=1}^N \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \wedge n\pi_i^p(\boldsymbol{\beta}_t) \wedge 1\} \\ & \quad \times \left| \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) - \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_t) \right| \|\mathbf{x}_i\|^2 \\ & \quad + \frac{1}{n} \sum_{i=1}^N |n\pi_i^p(\hat{\boldsymbol{\beta}}_1) - n\pi_i^p(\boldsymbol{\beta}_t)| \left| n\pi_i^p(\hat{\boldsymbol{\beta}}_1) + n\pi_i^p(\boldsymbol{\beta}_t) + 2 \right| \|\mathbf{x}_i\|^2 \\ & \equiv \Delta_6 + \Delta_7. \end{aligned} \quad (\text{A.62})$$

Note that $n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \wedge n\pi_i^p(\boldsymbol{\beta}_t) \wedge 1 \leq n\pi_i^p(\hat{\boldsymbol{\beta}}_1)$. Thus Δ_6 is bounded by

$$\frac{1}{\hat{\Psi}_1} \frac{1}{N} \sum_{i=1}^N \left| \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}_n) - \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \boldsymbol{\beta}_t) \right| h(\mathbf{x}_i) \|\mathbf{x}_i\|^2, \quad (\text{A.63})$$

which is $o_P(1)$ by Lemma 1 if $|\{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} - \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\}| = o_P(1)$. This is true because

$$\begin{aligned} |\{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} - \{n\pi_i^p(\boldsymbol{\beta}_t) \vee 1\}| & \leq n|\pi_i^p(\hat{\boldsymbol{\beta}}_1) - \pi_i^p(\boldsymbol{\beta}_t)| \\ & \leq \frac{nh(\mathbf{x}_i)}{N} \left| \frac{|\psi_i(\hat{\boldsymbol{\beta}}_1)|}{\hat{\Psi}_1} - \frac{|\psi_i(\boldsymbol{\beta}_t)|}{\Psi_N(\boldsymbol{\beta}_t)} \right| = o_P(1). \end{aligned}$$

The term Δ_7 is bounded by

$$\frac{1}{N} \sum_{i=1}^N \left| \frac{|\psi_i(\hat{\boldsymbol{\beta}}_1)|}{\hat{\Psi}_1} - \frac{|\psi_i(\boldsymbol{\beta}_t)|}{\Psi_N(\boldsymbol{\beta}_t)} \right| \left| \frac{|\psi_i(\hat{\boldsymbol{\beta}}_1)|}{\hat{\Psi}_1} + \frac{|\psi_i(\boldsymbol{\beta}_t)|}{\Psi_N(\boldsymbol{\beta}_t)} + \frac{2}{h(\mathbf{x}_i)} \right| h^2(\mathbf{x}_i) \|\mathbf{x}_i\|^2 = o_P(1),$$

where the equality to $o_P(1)$ is from Lemma 1 and the fact that $\mathbb{E}\{h^2(\mathbf{x})\|\mathbf{x}\|^2\} < \infty$. \square

A.3 Proofs for unconditional distribution

In this section we prove Theorem 3 in Section 6. A lemma similar to Lemma 5 is presented below. Lemma 6 can be used in the proof of Theorem 2 because for the problem considered in this paper, convergence to zero in probability is equivalent to convergence to zero in probability under the conditional probability measure (Xiong and Li, 2008).

For the pilot sample taken according to the subsampling probabilities π_{1i} in (25), we define $\delta_i^{(1)} = I\{u_{1i} \leq \frac{c_0(1-y_i)+c_1y_i}{N}\}$, where u_{1i} are i.i.d. standard uniform random variables. With this notation, the estimator $\hat{\Psi}_1$ defined in (26) can be written as

$$\hat{\Psi}_1 = \frac{1}{N} \sum_{i=1}^N \frac{\delta_i^{(1)} |y_i - p_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i)}{n\pi_{1i} \wedge 1}. \quad (\text{A.64})$$

Lemma 7. *Let $\hat{\boldsymbol{\beta}}_1$ and $\hat{\Psi}_1$ be constructed according to Step 1 of Algorithm 3, respectively. For*

$$\dot{\lambda}_p(\boldsymbol{\beta}_t) = \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \{y_i - p_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1)\} \mathbf{x}_i, \quad (\text{A.65})$$

under the same assumptions of Theorem 3, if $n = o(N)$, then

$$\frac{\dot{\lambda}_p(\boldsymbol{\beta}_t)}{\sqrt{n}} \longrightarrow \mathbb{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}), \quad (\text{A.66})$$

in distribution; if $n/N \rightarrow \rho \in (0, 1)$, then

$$\frac{\dot{\lambda}_p(\boldsymbol{\beta}_t)}{\sqrt{n}} \longrightarrow \mathbb{N}(\mathbf{0}, \boldsymbol{\Lambda}_u), \quad (\text{A.67})$$

in distribution.

Proof of Theorem 3. The proof of this theorem is similar to that of Theorem 2. The key difference is that Lemma 7 is about asymptotic distribution unconditionally.

The estimator $\hat{\boldsymbol{\beta}}_p$ is the maximizer of

$$\lambda_p(\boldsymbol{\beta}) = \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} [(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^T \mathbf{x}_i y_i - \log\{1 + e^{(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_1)^T \mathbf{x}_i}\}], \quad (\text{A.68})$$

so $\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t)$ is the maximizer of $\gamma_p(\mathbf{s}) = \lambda_p(\boldsymbol{\beta}_t + \mathbf{s}/\sqrt{n}) - \lambda_p(\boldsymbol{\beta}_t)$. By Taylor's expansion,

$$\gamma_p(\mathbf{s}) = \frac{1}{\sqrt{n}} \mathbf{s}^\top \dot{\lambda}_p(\boldsymbol{\beta}_t) + \frac{1}{2n} \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \mathbf{s}/\sqrt{n}) (\mathbf{s}^\top \mathbf{x}_i)^2 \quad (\text{A.69})$$

where $\acute{\mathbf{s}}$ lies between $\mathbf{0}$ and \mathbf{s} .

From Lemma 5,

$$\frac{1}{n} \sum_{i=1}^N \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \phi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1 + \acute{\mathbf{s}}/\sqrt{n}) \mathbf{x}_i (\mathbf{x}_i)^\top = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} + o_P(1). \quad (\text{A.70})$$

In addition, from Lemma 7, $\dot{\lambda}_p(\boldsymbol{\beta}_t)/\sqrt{n}$ converges in distribution to a normal limit. Thus, from the Basic Corollary in page 2 of Hjort and Pollard (2011), the maximizer of $\gamma_p(\mathbf{s})$, $\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t)$, satisfies

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_p - \boldsymbol{\beta}_t) = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t} \frac{1}{\sqrt{n}} \dot{\lambda}_p(\boldsymbol{\beta}_t) + o_P(1). \quad (\text{A.71})$$

Combining this with Lemma 7 and Slutsky's theorem, Theorem 3 follows. \square

A.3.1 Proof of Proposition 3

Proof of Proposition 3. To prove (39), we just need to show that $\boldsymbol{\Lambda}_u \geq \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} > \boldsymbol{\Lambda}_\rho$. From Proposition 2, we know that $\boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1} > \boldsymbol{\Lambda}_\rho$. To show that $\boldsymbol{\Lambda}_u \geq \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}$, we notice that

$$\begin{aligned} \boldsymbol{\Lambda}_u &= \frac{\mathbb{E}[\phi(\boldsymbol{\beta}_t) \{ \rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \vee \Phi(\boldsymbol{\beta}_t) \} h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top]}{4\Phi^2(\boldsymbol{\beta}_t)} \\ &\geq \frac{\mathbb{E}\{\phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \Phi(\boldsymbol{\beta}_t) \mathbf{x} \mathbf{x}^\top\}}{4\Phi^2(\boldsymbol{\beta}_t)} = \frac{\mathbb{E}\{\phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^\top\}}{4\Phi(\boldsymbol{\beta}_t)} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_t}^{-1}, \end{aligned}$$

where the strict inequality holds if $\rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \vee \Phi(\boldsymbol{\beta}_t) \neq \rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x})$ with positive probability, i.e., $\mathbb{P}\{\rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x}) > \Phi(\boldsymbol{\beta}_t)\} > 0$. \square

A.3.2 Proof of Lemma 7

Proof of Lemma 7. We first proof the case when the pilot estimates $\hat{\boldsymbol{\beta}}_1$ and $\hat{\Psi}_1$ depend on the data. For any $l \in \mathbb{R}^d$, denote $\tau_{N_i} = \sqrt{N/n} \hat{\Psi}_1 \delta_i^{\hat{\boldsymbol{\beta}}_1} \{n\pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \psi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i^\top l$, $i = 1, \dots, N$, where $\delta_i^{\hat{\boldsymbol{\beta}}_1} = I\{u_i \leq n\pi_i^p(\hat{\boldsymbol{\beta}}_1)\}$, and u_i are i.i.d. standard uniform random variables. Note that τ_{N_i} 's have the same distribution but they are not independent. We now exam the mean and variance of τ_{N_i} . For the mean, based on calculation similar to that in (A.22), we have,

$$\mathbb{E}(\tau_{N_i} | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) = \sqrt{nN} \hat{\Psi}_1 \mathbb{E}\{\pi_i^p(\hat{\boldsymbol{\beta}}_1) \psi_i(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) \mathbf{x}_i^\top l | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1\} = \frac{\sqrt{n} \mathbb{E}(\eta_i | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1)}{\sqrt{N}} = 0, \quad (\text{A.72})$$

which implies that

$$\mathbb{E}\tau_{Ni} = 0. \quad (\text{A.73})$$

For the variance, $\mathbb{V}(\tau_{Ni}) = \mathbb{E}(\tau_{Ni}^2)$, we start with the condition expectation,

$$\begin{aligned} \mathbb{E}(\tau_{Ni}^2 | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) &= N \hat{\Psi}_1^2 \mathbb{E} \left[\pi_i^p(\hat{\boldsymbol{\beta}}_1) \{n \pi_i^p(\hat{\boldsymbol{\beta}}_1) \vee 1\} \psi_i^2(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) (\mathbf{x}_i^T l)^2 \middle| \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1 \right] \\ &= \mathbb{E} \left[|\psi_i(\hat{\boldsymbol{\beta}}_1)| \left\{ \frac{n}{N} |\psi_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) \vee \hat{\Psi}_1 \right\} \psi_i^2(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 \middle| \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1 \right]. \end{aligned}$$

If we let

$$\Upsilon_{Ni} = |\psi_i(\hat{\boldsymbol{\beta}}_1)| \left\{ \frac{n}{N} |\psi_i(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_i) \vee \hat{\Psi}_1 \right\} \psi_i^2(\boldsymbol{\beta}_t - \hat{\boldsymbol{\beta}}_1) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2, \quad (\text{A.74})$$

then $\mathbb{V}(\tau_{Ni}) = \mathbb{E}(\Upsilon_{Ni})$. Note that

$$\Upsilon_{Ni} \rightarrow \Upsilon_i = 0.25 |\psi_i(\boldsymbol{\beta}_t)| \{ \rho_1 |\psi_i(\boldsymbol{\beta}_t)| h(\mathbf{x}_i) \vee \Psi(\boldsymbol{\beta}_t) \} h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2, \quad (\text{A.75})$$

in probability. We now show that

$$\mathbb{E}(\Upsilon_{Ni}) \rightarrow \mathbb{E}(\Upsilon_i) = \mathbb{E}[\phi(\boldsymbol{\beta}_t) \{ \rho \phi(\boldsymbol{\beta}_t) h(\mathbf{x}) \vee \Phi \phi(\boldsymbol{\beta}_t) \} h(\mathbf{x}) (\mathbf{x}^T l)^2]. \quad (\text{A.76})$$

Let $\Xi = |\Upsilon_{Ni} - \Upsilon_i|$. For any ϵ ,

$$\begin{aligned} |\mathbb{E}(\Upsilon_{Ni}) - \mathbb{E}(\Upsilon_i)| &= \mathbb{E}\{\Xi I(\Xi > \epsilon)\} + \mathbb{E}\{\Xi I(\Xi \leq \epsilon)\} \\ &\leq \mathbb{E}[\{h^2(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 + \Upsilon_i + \hat{\Psi}_1 h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2\} I(\Xi > \epsilon)] + \epsilon \end{aligned}$$

We know that $\mathbb{E}[\{h^2(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 + \Upsilon_i\} I(\Xi > \epsilon)] \rightarrow 0$ since $\mathbb{E}\{h^2(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 + \Upsilon_i\} < \infty$ for any $l \in \mathbb{R}^d$, and $I(\Xi > \epsilon)$ is bounded and is $o_P(1)$. Thus, to prove that $\mathbb{E}(\Upsilon_{Ni}) - \mathbb{E}(\Upsilon_i) \rightarrow 0$, we only need to show that

$$\begin{aligned} &\mathbb{E}\{\hat{\Psi}_1 h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 I(\Xi > \epsilon)\} \\ &= \frac{1}{N} \sum_{k=1}^N \mathbb{E} \left\{ \frac{\delta_k^{(1)} |y_k - p_k(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_k)}{n_1 \pi_k^{(1)} \wedge 1} h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 I(\Xi > \epsilon) \right\} \equiv \frac{1}{N} \sum_{k=1}^N \Delta_{8k} \rightarrow 0, \end{aligned}$$

which is true if $\Delta_{8k} = o(1)$ for any k . We need to keep in mind that $\delta_k^{(1)} = I\{u_{1k} \leq \frac{c_0(1-y_k) + c_1 y_k}{N}\}$ and Υ_{Ni} are correlated, and they both depend on the data. For π_{1k} with the expression in (25), we have

$$\frac{1}{n_1 \pi_{1k} \wedge 1} \leq \frac{N}{n_1} \frac{c_0 + c_1}{c_0 c_1} + 1 \quad \text{and} \quad n_1 \pi_{1k} \wedge 1 \leq \frac{n_1}{N} (c_0 + c_1). \quad (\text{A.77})$$

Using this, we have

$$\Delta_{8k} \leq \left(\frac{N}{n_1} \frac{c_0 + c_1}{c_0 c_1} + 1 \right) \mathbb{E} \left\{ \delta_k^{(1)} h(\mathbf{x}_k) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 I(\Xi > \epsilon) \right\} \quad (\text{A.78})$$

$$= \left(\frac{N}{n_1} \frac{c_0 + c_1}{c_0 c_1} + 1 \right) \mathbb{E} \left\{ h(\mathbf{x}_k) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 I(\Xi > \epsilon) \mid \delta_k^{(1)} = 1 \right\} \mathbb{P}(\delta_j^{(1)} = 1) \quad (\text{A.79})$$

$$< \frac{(c_0 + c_1 + c_0 c_1)^2}{c_0 c_1} \mathbb{E} \left\{ h(\mathbf{x}_k) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 I(\Xi > \epsilon) \mid \delta_k^{(1)} = 1 \right\} \quad (\text{A.80})$$

Note that $\mathbb{E} \{ h(\mathbf{x}_k) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^2 \mid \delta_k^{(1)} = 1 \} < \infty$ for any i, k , because $\mathbb{E} \{ h^2(\mathbf{x}) \|x\|^2 \} < \infty$. Thus, $\Delta_{8k} = o(1)$, if $\Xi = o_P(1)$ given $\delta_k^{(1)} = 1$. From the expression of Ξ , we only need to show that $\hat{\beta}_1$ and $\hat{\Psi}_1$ are consistent given $\delta_k^{(1)} = 1$. This is true because $n_1 \rightarrow \infty$ and fixing the value of one observation does not affect the consistency. For example, for $\hat{\Psi}_1$, given $\delta_k^{(1)} = 1$,

$$\hat{\Psi}_1 = \frac{1}{N} \sum_{j \neq k}^N \frac{\delta_j^{(1)} |y_j - p_j(\hat{\beta}_1)| h(\mathbf{x}_j)}{n_1 \pi_j^{(1)} \wedge 1} + \frac{|y_k - p_k(\hat{\beta}_1)| h(\mathbf{x}_k)}{n_1 \{c_0(1 - y_k) + c_1 y_k\} \wedge N}, \quad (\text{A.81})$$

in which the first term converge to $\Psi(\beta_t) = 2\Phi(\beta_t)$ and the second term is $o_P(1)$. We have finished proving that

$$\mathbb{V}(\tau_{Ni}) \rightarrow \mathbb{E}(\Upsilon_i) = \mathbb{E}[\phi(\beta_t) \{ \rho \phi(\beta_t) h(\mathbf{x}) \vee \Phi \phi(\beta_t) \} h(\mathbf{x}) (\mathbf{x}^T l)^2]. \quad (\text{A.82})$$

In the following, we exam the third moment of τ_{Ni} and prove that

$$\mathbb{E} |\tau_{Ni}|^3 = o(\sqrt{N}). \quad (\text{A.83})$$

For the conditional expectation,

$$\begin{aligned} & \mathbb{E} (|\tau_{Ni}|^3 \mid \hat{\beta}_1, \hat{\Psi}_1) \\ &= N \sqrt{N/n} \hat{\Psi}_1^3 \mathbb{E} \left[\pi_i^p(\hat{\beta}_1) \{ n \pi_i^p(\hat{\beta}_1) \vee 1 \}^2 \psi_i^3(\beta_t - \hat{\beta}_1) (\mathbf{x}_i^T l)^3 \mid \hat{\beta}_1, \hat{\Psi}_1 \right] \\ &= \sqrt{N/n} \mathbb{E} \left[|\psi_i(\hat{\beta}_1)| \left\{ \frac{n}{N} |\psi_i(\hat{\beta}_1)| h(\mathbf{x}_i) \vee \hat{\Psi}_1 \right\}^2 \psi_i^3(\beta_t - \hat{\beta}_1) h(\mathbf{x}_i) (\mathbf{x}_i^T l)^3 \mid \hat{\beta}_1, \hat{\Psi}_1 \right] \\ &\leq 2 \|l\|^3 \sqrt{N/n} \mathbb{E} [\{h^2(\mathbf{x}_i) + \hat{\Psi}_1^2\} h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \mid \hat{\beta}_1, \hat{\Psi}_1]. \end{aligned}$$

Thus, (A.83) follows if $\mathbb{E} [\{h^2(\mathbf{x}_i) + \hat{\Psi}_1^2\} h(\mathbf{x}_i) \|\mathbf{x}_i\|^3] = O(1)$, which is true if

$$\mathbb{E} \{ \hat{\Psi}_1^2 h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \} = O(1) \quad (\text{A.84})$$

since $\mathbb{E} \{ h^3(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \} = \mathbb{E} \{ h^3(\mathbf{x}) \|\mathbf{x}\|^3 \} < \infty$ and it is a constant that does not depend n_1, n , or N . For (A.84),

$$\mathbb{E} \{ \hat{\Psi}_1^2 h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \}$$

$$\begin{aligned}
&= \mathbb{E} \left\{ \frac{1}{N} \sum_{k_1=1}^N \frac{\delta_{k_1}^{(1)} |y_{k_1} - p_{k_1}(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_{k_1})}{n_1 \pi_{k_1}^{(1)} \wedge 1} \frac{1}{N} \sum_{k_2=1}^N \frac{\delta_{k_2}^{(1)} |y_{k_2} - p_{k_2}(\hat{\boldsymbol{\beta}}_1)| h(\mathbf{x}_{k_2})}{n_1 \pi_{k_2}^{(1)} \wedge 1} h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \right\} \\
&\leq \frac{1}{N^2} \sum_{k_1 \neq k_2}^N \mathbb{E} \left\{ \frac{\delta_{k_1}^{(1)} h(\mathbf{x}_{k_1})}{n_1 \pi_{k_1}^{(1)} \wedge 1} \frac{\delta_{k_2}^{(1)} h(\mathbf{x}_{k_2})}{n_1 \pi_{k_2}^{(1)} \wedge 1} h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \right\} + \frac{1}{N^2} \sum_{k=1}^N \mathbb{E} \left\{ \frac{\delta_k^{(1)} h^2(\mathbf{x}_k)}{\{n_1 \pi_{1k} \wedge 1\}^2} h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \right\} \\
&= \frac{1}{N^2} \sum_{k_1 \neq k_2}^N \mathbb{E} \{ h(\mathbf{x}_{k_1}) h(\mathbf{x}_{k_2}) h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \} + \frac{1}{N^2} \sum_{j=1}^N \mathbb{E} \left\{ \frac{1}{n_1 \pi_{1k} \wedge 1} h^2(\mathbf{x}_k) h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \right\} \\
&\leq \frac{1}{N^2} \sum_{k_1 \neq k_2}^N \mathbb{E} \{ h(\mathbf{x}_{k_1}) h(\mathbf{x}_{k_2}) h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \} + \frac{c_0 + c_1 + c_0 c_1}{c_0 c_1 n_1 N} \sum_{j=1}^N \mathbb{E} \{ h^2(\mathbf{x}_k) h(\mathbf{x}_i) \|\mathbf{x}_i\|^3 \} \\
&= \frac{N-1}{N} \{ \mathbb{E} h(\mathbf{x}) \}^2 \mathbb{E} \{ h(\mathbf{x}) \|\mathbf{x}\|^3 \} + \frac{c_0 + c_1 + c_0 c_1}{c_0 c_1 n_1} [\mathbb{E} \{ h^2(\mathbf{x}) \}] [\mathbb{E} \{ h(\mathbf{x}) \|\mathbf{x}\|^3 \}] = O(1). \quad (\text{A.85})
\end{aligned}$$

This finish the proof of (A.84).

Denote $\nu_{Ni} = \tau_{Ni} \{ \mathbb{V}(\tau_{Ni}) \}^{-1/2}$. We know that ν_{Ni} are i.i.d. conditional on $\hat{\boldsymbol{\beta}}_1$ and $\hat{\Psi}_1$. Thus, from Theorem 7.3.2 of Chow and Teicher (2003), they are interchangeable. The fact that $\hat{\boldsymbol{\beta}}_1$ and $\hat{\Psi}_1$ are consistent estimators implies that they are a sequence of two estimators, and for each $\hat{\boldsymbol{\beta}}_1$ and $\hat{\Psi}_1$, τ_{Ni} are interchangeable. For this setup, the central limit theorem in Theorem 2 of Blum et al. (1958) can be applied to prove the asymptotic normality.

It is evident that ν_{Ni} have mean 0 and variance 1. It is also easy to verify that, for $i \neq j$,

$$\mathbb{E}(\nu_{Ni} \nu_{Nj}) = \mathbb{E} \{ \mathbb{E}(\nu_{Ni} \nu_{Nj} | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) \} = 0, \quad (\text{A.86})$$

and

$$\frac{1}{\sqrt{N}} \mathbb{E} \{ |\nu_{Ni}|^3 \} = \mathbb{E} |\tau_{Ni}|^3 \{ \mathbb{V}(\tau_{Ni}) \}^{-3/2} \rightarrow 0 \quad (\text{A.87})$$

which follows from (A.83). We now show that for $i \neq j$,

$$\mathbb{E} \{ \nu_{Ni}^2 \nu_{Nj}^2 \} \rightarrow 1. \quad (\text{A.88})$$

Since $\nu_{Ni} = \tau_{Ni} \{ \mathbb{V}(\tau_{Ni}) \}^{-1/2}$, from (A.82), to prove (A.88), we only need to show that $\mathbb{E}(\tau_{Ni}^2 \tau_{Nj}^2) \rightarrow \mathbb{E}(\Upsilon_i) \mathbb{E}(\Upsilon_j) = \mathbb{E}(\Upsilon_i \Upsilon_j)$, where the equality is because Υ_i and Υ_j are independent. Noting that τ_{Ni}^2 and τ_{Nj}^2 are conditionally independent, we have $\mathbb{E}(\tau_{Ni}^2 \tau_{Nj}^2 | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) = \mathbb{E}(\tau_{Ni}^2 | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) \mathbb{E}(\tau_{Nj}^2 | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) = \mathbb{E}(\Upsilon_{Ni} | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) \mathbb{E}(\Upsilon_{Nj} | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1) = \mathbb{E}(\Upsilon_{Ni} \Upsilon_{Nj} | \hat{\boldsymbol{\beta}}_1, \hat{\Psi}_1)$, so we know that $\mathbb{E}(\tau_{Ni}^2 \tau_{Nj}^2) = \mathbb{E}(\Upsilon_{Ni} \Upsilon_{Nj})$.

Now we prove that $\mathbb{E}(\Upsilon_{Ni} \Upsilon_{Nj}) \rightarrow \mathbb{E}(\Upsilon_i \Upsilon_j)$. Let $\Xi_2 = |\Upsilon_{Ni} \Upsilon_{Nj} - \Upsilon_i \Upsilon_j|$. For any $\epsilon > 0$,

$$\begin{aligned}
&|\mathbb{E}(\Upsilon_{Ni} \Upsilon_{Nj}) - \mathbb{E}(\Upsilon_i \Upsilon_j)| \\
&\leq \mathbb{E} \{ \Xi_2 I(\Xi_2 > \epsilon) \} + \mathbb{E} \{ \Xi_2 I(\Xi_2 \leq \epsilon) \}
\end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{E}\left[\{h^2(\mathbf{x}_i)(\mathbf{x}_i^\top l)^2 h^2(\mathbf{x}_j)(\mathbf{x}_j^\top l)^2 + \Upsilon_i \Upsilon_j\} I(\Xi_2 > \epsilon)\right] \\
&\quad + \mathbb{E}\left[\hat{\Psi}_1^2 h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2 I(\Xi_2 > \epsilon)\right] \\
&\quad + \mathbb{E}\left[\hat{\Psi}_1 \{h(\mathbf{x}_i)(\mathbf{x}_i^\top l)^2 h^2(\mathbf{x}_j)(\mathbf{x}_j^\top l)^2 + h(\mathbf{x}_j)(\mathbf{x}_j^\top l)^2 h^2(\mathbf{x}_i)(\mathbf{x}_i^\top l)^2\} I(\Xi_2 > \epsilon)\right] + \epsilon \\
&\equiv \Delta_9 + \Delta_{10} + \Delta_{11} \tag{A.89}
\end{aligned}$$

Note that $\mathbb{E}\{h^2(\mathbf{x}_i)(\mathbf{x}_i^\top l)^2 h^2(\mathbf{x}_j)(\mathbf{x}_j^\top l)^2 + \Upsilon_i \Upsilon_j\} < \infty$, $I(\Xi_2 > \epsilon) \leq 1$, and $\Xi_2 = o_P(1)$, thus $\Delta_9 \rightarrow 0$. Now we exam Δ_{10} and Δ_{11} show that they are also $o(1)$. For Δ_{10} ,

$$\begin{aligned}
\Delta_{10} &= \mathbb{E}\left\{\hat{\Psi}_1^2 h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2 I(\Xi_2 > \epsilon)\right\} \\
&= \mathbb{E}\left\{\frac{1}{N} \sum_{k_1=1}^N \frac{\delta_{k_1}^{(1)} |y_{k_1} - p_{k_1}(\hat{\beta}_1)| h(\mathbf{x}_{k_1})}{n_1 \pi_{k_1}^{(1)} \wedge 1} \right. \\
&\quad \left. \times \frac{1}{N} \sum_{k_2=1}^N \frac{\delta_{k_2}^{(1)} |y_{k_2} - p_{k_2}(\hat{\beta}_1)| h(\mathbf{x}_{k_2})}{n_1 \pi_{k_2}^{(1)} \wedge 1} h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2 I(\Xi_2 > \epsilon)\right\} \\
&\leq \frac{1}{N^2} \sum_{k_1 \neq k_2}^N \mathbb{E}\left\{\frac{\delta_{k_1}^{(1)} h(\mathbf{x}_{k_1})}{n_1 \pi_{k_1}^{(1)} \wedge 1} \frac{\delta_{k_2}^{(1)} h(\mathbf{x}_{k_2})}{n_1 \pi_{k_2}^{(1)} \wedge 1} h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2 I(\Xi_2 > \epsilon)\right\} \tag{A.90}
\end{aligned}$$

$$+ \frac{1}{N^2} \sum_{k=1}^N \mathbb{E}\left\{\frac{\delta_k^{(1)} h^2(\mathbf{x}_k)}{\{n_1 \pi_{1k} \wedge 1\}^2} h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2 I(\Xi_2 > \epsilon)\right\}. \tag{A.91}$$

Let $h_{k_1 k_2 i j} x_{i^2 j^2} = h(\mathbf{x}_{k_1}) h(\mathbf{x}_{k_2}) h(\mathbf{x}_i) \|\mathbf{x}_i\|^2 h(\mathbf{x}_j) \|\mathbf{x}_j\|^2$. From (A.77), we have

$$\begin{aligned}
&\mathbb{E}\left\{\frac{\delta_{k_1}^{(1)}}{n_1 \pi_{k_1}^{(1)} \wedge 1} \frac{\delta_{k_2}^{(1)}}{n_1 \pi_{k_2}^{(1)} \wedge 1} h_{k_1 k_2 i j} x_{i^2 j^2} I(\Xi_2 > \epsilon)\right\} \\
&\leq \left(\frac{N}{n_1} \frac{c_0 + c_1}{c_0 c_1} + 1\right)^2 \mathbb{E}\left[h_{k_1 k_2 i j} x_{i^2 j^2} I(\Xi_2 > \epsilon) \Big| \delta_{k_1}^{(1)} = \delta_{k_2}^{(1)} = 1\right] \mathbb{P}(\delta_{k_1}^{(1)} = \delta_{k_2}^{(1)} = 1) \\
&< \frac{(c_0 + c_1 + c_0 c_1)^2 (c_0 + c_1)^2}{c_0^2 c_1^2} \mathbb{E}\left[h_{k_1 k_2 i j} x_{i^2 j^2} I(\Xi_2 > \epsilon) \Big| \delta_{k_1}^{(1)} = \delta_{k_2}^{(1)} = 1\right]. \tag{A.92}
\end{aligned}$$

Note that although k_1 or k_2 may equal i or j , $k_1 \neq k_2$, so we know that $\mathbb{E}(h_{k_1 k_2 i j} x_{i^2 j^2}) < \infty$ since $\mathbb{E}\{h^2(\mathbf{x}) \|\mathbf{x}\|^2\} < \infty$. Thus, the term in (A.92) is $o(1)$ if $\Xi_2 = o_P(1)$ given $\delta_j^{(1)} = 1$. This is true because $\hat{\beta}_1$ and $\hat{\Psi}_1$ are consistent given $\delta_{k_1}^{(1)} = \delta_{k_2}^{(1)} = 1$, namely, fixing the values of two observations does not affect the consistency of $\hat{\beta}_1$ and $\hat{\Psi}_1$. From (A.92), the term in (A.90) converges to zero. Using a similar approach, it can be shown that if $\mathbb{E}\{h^3(\mathbf{x}) \|\mathbf{x}\|^2\} < \infty$, then $\Delta_{11} = o(1)$, and the term in (A.91) converges to zero, which implies that $\Delta_{10} = o(1)$. Thus, (A.88) holds.

Since (A.86), (A.87), (A.88) are satisfied, the central limit theorem in Theorem 2 of Blum et al. (1958) holds for ν_{Ni} , which gives that

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N \nu_{Ni} \rightarrow \mathbb{N}(0, 1), \tag{A.93}$$

in distribution. Note that

$$\begin{aligned} \frac{1}{\sqrt{N}} \sum_{i=1}^N \nu_i &= \frac{\hat{\Psi}_1}{\sqrt{n}\{\mathbb{V}(\tau_{Ni})\}^{1/2}} \sum_{i=1}^N \delta_i^{\hat{\beta}_1} \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i(\beta_t - \hat{\beta}_1) \mathbf{x}_i^T l \\ &= \frac{\hat{\Psi}_1}{\sqrt{n}\{\mathbb{V}(\tau_{Ni})\}^{1/2}} l^T \dot{\lambda}_p(\beta_t) = \frac{\Psi}{\sqrt{n}\{\mathbb{V}(\tau_{Ni})\}^{1/2}} l^T \dot{\lambda}_p(\beta_t) + o_P(1). \end{aligned}$$

Thus, from Slutsky's theorem, for any $l \in \mathbb{R}^d$,

$$\frac{1}{\sqrt{n}} l^T \dot{\lambda}_p(\beta_t) \rightarrow \mathbb{N}(0, l^T \mathbf{\Lambda}_u l) \quad (\text{A.94})$$

in distribution, where

$$\mathbf{\Lambda}_u = \frac{\mathbb{V}(\tau_{Ni})}{\Psi^2(\beta_t)} = \frac{\mathbb{E}[\phi(\beta_t) \{ \rho \phi(\beta_t) h(\mathbf{x}) \vee \Phi \} h(\mathbf{x}) \mathbf{x} \mathbf{x}^T]}{4\Phi^2(\beta_t)} \geq \frac{\mathbb{E}[\phi(\beta_t) h(\mathbf{x}) \mathbf{x} \mathbf{x}^T]}{4\Phi(\beta_t)} = \mathbf{\Sigma}_{\beta_t}^{-1}, \quad (\text{A.95})$$

and the equality holds if $\rho = 0$, i.e., $n/N \rightarrow 0$. Based on (A.94), from the Cramér-Wold theorem, we have that

$$\frac{1}{\sqrt{n}} \dot{\lambda}_p(\beta_t) \rightarrow \mathbb{N}(0, \mathbf{\Lambda}_u) \quad (\text{A.96})$$

in distribution.

When the pilot estimates $\hat{\beta}_1$ and $\hat{\Psi}_1$ are independent of the data, if we can prove the results in Lemma 7 under the conditional distribution given $\hat{\beta}_1$ and $\hat{\Psi}_1$, then the result follows unconditionally. We provide the proof under the conditional distribution in the following. The proof is similar to the proof of Lemma 5 and thus we provide only the outline. The difference is we do not conditional on the full data \mathcal{D}_N here.

Note that, given $\hat{\beta}_1$ and $\hat{\Psi}_1$, $\dot{\lambda}_p(\beta_t)$ is a sum of N independent random vectors. We now exam the mean and variance of $\dot{\lambda}_p(\beta_t)$ given $\hat{\beta}_1$ and $\hat{\Psi}_1$. For the mean,

$$\frac{1}{\sqrt{n}} \mathbb{E}\{\dot{\lambda}_p(\beta_t) | \hat{\beta}_1, \hat{\Psi}_1\} = \mathbf{0}. \quad (\text{A.97})$$

For the variance,

$$\frac{1}{n} \mathbb{V}\{\dot{\lambda}_p(\beta_t) | \hat{\beta}_1, \hat{\Psi}_1\} = \frac{\mathbb{E}\left[|\psi_i(\hat{\beta}_1)| \{n\pi_i^p(\hat{\beta}_1) \vee 1\} \psi_i^2(\beta_t - \hat{\beta}_1) h(\mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^T \middle| \hat{\beta}_1, \hat{\Psi}_1\right]}{\hat{\Psi}_1}, \quad (\text{A.98})$$

which, under Assumptions 1 and 2, converges in probability to $\mathbf{\Lambda}_u$.

To check the Lindeberg-Feller condition (Section *2.8 of van der Vaart, 1998) under the condition distribution, we note that for any $\epsilon > 0$,

$$\frac{1}{n} \sum_{i=1}^N \mathbb{E}\left\{ \|\dot{\lambda}_{pi}\|^2 I(\|\dot{\lambda}_{pi}\| > \sqrt{n}\epsilon) \middle| \hat{\beta}_1, \hat{\Psi}_1 \right\}$$

$$\leq \frac{\mathbb{E} \left[\{h(\mathbf{x})\|\mathbf{x}\|^2 + \hat{\Psi}_1\} h(\mathbf{x}) I(\{h(\mathbf{x})/\hat{\Psi}_1 + 1\}\|\mathbf{x}\| > \sqrt{n}\epsilon) \middle| \hat{\beta}_1, \hat{\Psi}_1 \right]}{\hat{\Psi}_1^2} = o_P(1).$$

Thus, applying the Lindeberg-Feller central limit theorem (Section *2.8 of van der Vaart, 1998) finishes the proof. □

References

- Avron, H., Maymounkov, P. and Toledo, S. (2010), ‘Blendenpik: Supercharging LAPACK’s least-squares solver’, *SIAM Journal on Scientific Computing* **32**, 1217–1236.
- Baldi, P., Sadowski, P. and Whiteson, D. (2014), ‘Searching for exotic particles in high-energy physics with deep learning’, *Nature Communications* **5**(4308), <http://dx.doi.org/10.1038/ncomms5308>.
- Bezanson, J., Edelman, A., Karpinski, S. and Shah, V. B. (2017), ‘Julia: A fresh approach to numerical computing’, *SIAM Review* **59**(1), 65–98.
- Blum, J., Chernoff, H., Rosenblatt, M. and Teicher, H. (1958), ‘Central limit theorems for interchangeable processes’, *Canad. J. Math* **10**, 222–229.
- Chow, Y. S. C. and Teicher, H. (2003), *Probability Theory: Independence, Interchangeability, Martingales*, Springer, New York.
- Clarkson, K. L. and Woodruff, D. P. (2013), Low rank approximation and regression in input sparsity time, *in* ‘Proceedings of the forty-fifth annual ACM symposium on Theory of computing’, ACM, pp. 81–90.
- Drineas, P., Kannan, R. and Mahoney, M. W. (2006a), ‘Fast monte carlo algorithms for matrices i: Approximating matrix multiplication’, *SIAM Journal on Computing* **36**(1), 132–157.
- Drineas, P., Kannan, R. and Mahoney, M. W. (2006b), ‘Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix’, *SIAM Journal on Computing* **36**(1), 158–183.
- Drineas, P., Kannan, R. and Mahoney, M. W. (2006c), ‘Fast monte carlo algorithms for matrices iii: Computing a compressed approximate matrix decomposition’, *SIAM Journal on Computing* **36**(1), 184–206.

- Drineas, P., Magdon-Ismail, M., Mahoney, M. and Woodruff, D. (2012), ‘Faster approximation of matrix coherence and statistical leverage.’, *Journal of Machine Learning Research* **13**, 3475–3506.
- Drineas, P., Mahoney, M., Muthukrishnan, S. and Sarlos, T. (2011), ‘Faster least squares approximation’, *Numerische Mathematik* **117**, 219–249.
- Drineas, P., Mahoney, M. W. and Muthukrishnan, S. (2006), Sampling algorithms for l_2 regression and applications, *in* ‘Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm’, Society for Industrial and Applied Mathematics, pp. 1127–1136.
- Fithian, W. and Hastie, T. (2014), ‘Local case-control sampling: Efficient subsampling in imbalanced data sets’, *Annals of statistics* **42**(5), 1693.
- Halko, N., Martinsson, P.-G. and Tropp, J. A. (2011), ‘Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions’, *SIAM review* **53**(2), 217–288.
- Hjort, N. L. and Pollard, D. (2011), ‘Asymptotics for minimisers of convex processes’, *arXiv preprint arXiv:1107.3806*.
- Hosmer Jr, D. W., Lemeshow, S. and Sturdivant, R. X. (2013), *Applied logistic regression*, Vol. 398, John Wiley & Sons.
- Kernighan, B. W. and Ritchie, D. M. (1988), *The C Programming Language*, 2nd edn, Prentice Hall Professional Technical Reference.
- Kleiner, A., Talwalkar, A., Sarkar, P. and Jordan, M. I. (2014), ‘A scalable bootstrap for massive data’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(4), 795–816.
- Lichman, M. (2013), ‘UCI Machine Learning Repository [<http://archive.ics.uci.edu/ml>]. Irvine, CA: University of California, School of Information and Computer Science’.
URL: <http://archive.ics.uci.edu/ml>
- Ma, P., Mahoney, M. and Yu, B. (2014), A statistical perspective on algorithmic leveraging, *in* ‘Proceedings of the 31st International Conference on Machine Learning (ICML-14)’, pp. 91–99.

- Ma, P., Mahoney, M. and Yu, B. (2015), ‘A statistical perspective on algorithmic leveraging’, *Journal of Machine Learning Research* **16**, 861–911.
- Ma, P. and Sun, X. (2015), ‘Leveraging for big data regression’, *Wiley Interdisciplinary Reviews: Computational Statistics* **7**(1), 70–76.
URL: <http://dx.doi.org/10.1002/wics.1324>
- Mahoney, M. W. (2011), ‘Randomized algorithms for matrices and data’, *Foundations and Trends® in Machine Learning* **3**(2), 123–224.
- Mahoney, M. W. and Drineas, P. (2009), ‘CUR matrix decompositions for improved data analysis’, *Proceedings of the National Academy of Sciences* **106**(3), 697–702.
- McCullagh, P. and Nelder, J. A. (1989), *Generalized Linear Models, no. 37 in Monograph on Statistics and Applied Probability*, Chapman & Hall,.
- McWilliams, B., Krummenacher, G., Lucic, M. and Buhmann, J. M. (2014), Fast and robust least squares estimation in corrupted linear models, *in* ‘Advances in Neural Information Processing Systems’, pp. 415–423.
- Meng, X., Saunders, M. and Mahoney, M. (2014), ‘LSRN: A parallel iterative solver for strongly over- or under- determined systems’, *SIAM Journal on Scientific Computing* **36**, C95–C118.
- R Core Team (2017), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: <https://www.R-project.org/>
- Raskutti, G. and Mahoney, M. (2016), ‘A statistical perspective on randomized sketching for ordinary least-squares’, *Journal of Machine Learning Research* **17**, 1–31.
- Stroustrup, B. (1986), *The C++ programming language*, Pearson Education India.
- van der Vaart, A. (1998), *Asymptotic Statistics*, Cambridge University Press, London.
- van Rossum, G. (1995), Python tutorial, technical report cs-r9526, Technical report, Centrum voor Wiskunde en Informatica (CWI), Amsterdam.
- Wang, H., Yang, M. and Stufken, J. (2018), ‘Information-based optimal subdata selection for big data linear regression’, *Journal of the American Statistical Association* **accepted**, <https://doi.org/10.1080/01621459.2017.1408468>.

- Wang, H., Zhu, R. and Ma, P. (2017), ‘Optimal subsampling for large sample logistic regression’, *Journal of the American Statistical Association* p. <http://dx.doi.org/10.1080/01621459.2017.1292914>.
- Xiong, S. and Li, G. (2008), ‘Some results on the convergence of conditional distributions’, *Statistics & Probability Letters* **78**(18), 3249–3253.
- Yang, T., Zhang, L., Jin, R. and Zhu, S. (2015), An explicit sampling dependent spectral error bound for column subset selection, *in* ‘Proceedings of The 32nd International Conference on Machine Learning’, pp. 135–143.
- Yang, Y., Pilanci, M. and Wainwright, M. J. (2016), ‘Randomized sketches for kernels: Fast and optimal non-parametric regression’, *The Annals of Statistics*, p. forthcoming.