

Inference for Deep Neural Network Estimators in Generalized Nonparametric Models

Xuran Meng* and Yi Li†

Abstract

While deep neural networks (DNNs) are used for prediction, inference on DNN-estimated subject-specific means for categorical or exponential family outcomes remains underexplored. We address this by proposing a DNN estimator under generalized nonparametric regression models (GNRMs) and developing a rigorous inference framework. Unlike existing approaches that assume independence between estimation errors and inputs to establish the error bound, a condition often violated in GNRMs, we allow for dependence and our theoretical analysis demonstrates the feasibility of drawing inference under GNRMs. To implement inference, we consider an Ensemble Subsampling Method (ESM) that leverages U-statistics and the Hoeffding decomposition to construct reliable confidence intervals for DNN estimates. We show that, under GNRM settings, ESM enables model-free variance estimation and accounts for heterogeneity among individuals in the population. Through simulations under nonparametric logistic, Poisson, and binomial regression models, we demonstrate the effectiveness and efficiency of our method. We further apply the method to the electronic Intensive Care Unit (eICU) dataset, a large scale collection of anonymized health records from ICU patients, to predict ICU readmission risk and offer patient-centric insights for clinical decision making.

Keywords: Deep Neural Networks, Inference, Generalized Nonparametric Regression Models, Ensemble Subsampling Method

*Department of Biostatistics, University of Michigan; e-mail: xuranm@umich.edu

†Department of Biostatistics, University of Michigan; e-mail: yili@umich.edu

1 Introduction

Given an outcome variable y and covariates \mathbf{x} , estimating the conditional expectation $\mathbb{E}(y|\mathbf{x})$ is central to supervised learning. Under squared loss, $\mathbb{E}(y|\mathbf{x})$ is the optimal estimator that minimizes the expected estimation error among all measurable functions, making it a natural and widely adopted target in statistical modeling and machine learning. Even when optimality does not hold under more general losses, $\mathbb{E}(y|\mathbf{x})$ remains a valuable target, offering insight into the relationship between covariates and outcomes. It serves as a foundational component for constructing more complex quantities across domains such as conditional variances, counterfactual means, and risk-adjusted estimates (Hastie et al. 2009). Methods that estimate $\mathbb{E}(y|\mathbf{x})$ include linear, parametric and nonparametric regression, kernel regression (Nadaraya 1964), support vector regression (Drucker et al. 1996), random forests (Breiman 2001) and deep neural networks (DNNs) (LeCun et al. 2015).

With the rapid advancement of artificial intelligence, DNN-based estimators have gained widespread popularity. However, the statistical inference surrounding the DNN estimator of $\mathbb{E}(y|\mathbf{x})$, especially when y is categorical or belongs to the exponential family, remains underdeveloped. We address this gap by establishing a rigorous framework for inference on $\mathbb{E}(y|\mathbf{x})$ estimated by DNNs under a generalized nonparametric regression model (GNRM) setting [see Eq. (2.3)].

Modern inference approaches include distribution-free conformal estimation inference (Lei et al. 2018, Angelopoulos & Bates 2021, Huang et al. 2024), extensions of jackknife inference within conformal estimation (Alaa & Van Der Schaar 2020, Kim et al. 2020, Barber et al. 2021), debiasing methods (Athey et al. 2018, Guo et al. 2021), residual-based bootstrap techniques (Zhang & Politis 2023), and the use of U statistics in regressions, random forests and neural networks (Mentch & Hooker 2016, Wager & Athey 2018, Schupbach

et al. 2020, Wang & Wei 2022, Fei & Li 2024). Several distinctions and limitations remain. For example, conformal prediction constructs distribution-free prediction sets for y , targeting outcome coverage rather than inference on $\mathbb{E}(y|\mathbf{x})$. Its intervals thus absorb both outcome noise and model uncertainty, making them conservative and inefficient for mean inference. Clinically, inference on risk classifications derived from patients’ risk profiles, which are based on the estimated conditional mean (i.e., risk score), is more actionable, as medical decisions (e.g., triage or treatment intensity) rely on risk stratifications rather than prediction sets (Lundberg et al. 2018, Van Calster et al. 2019). Jackknife and related leave-one-out methods assume stable, nearly independent residuals, conditions violated by DNNs whose estimates depend heavily on individual samples. Debiasing approaches (Athey et al. 2018, Guo et al. 2021) and residual bootstraps (Zhang & Politis 2023) remain limited to linear models, while resampling-based inference for high-dimensional GLMs (Fei & Li 2021) is fully parametric. Recent work on uncertainty quantification for random forests (Mentch & Hooker 2016, Wager & Athey 2018) and neural networks (Schupbach et al. 2020, Fei & Li 2024) offers valuable insights but remains largely confined to continuous outcomes, leaving inference for GNRMs such as logistic and Poisson regressions unresolved.

Theoretically, for valid inference on $\mathbb{E}(y|\mathbf{x})$, we need to establish the convergence rate of neural networks in a GNRM framework. Since the universal approximation theorem (Hornik et al. 1989), there has been much progress in understanding the convergence rates of neural networks (McCaffrey & Gallant 1994, Kohler & Krzyżak 2005, Yarotsky 2017, Shen et al. 2019, Kohler & Langer 2021, Lu et al. 2021, Shen et al. 2021, Jiao et al. 2023, Fan et al. 2024, Wang et al. 2024, Bhattacharya et al. 2024, Yan et al. 2025). For instance, Shen et al. (2021) established the convergence rate under nonparametric regression assuming the p -th moment of response is bounded for some $p > 1$; Fan et al. (2024)

explored nonasymptotic error bounds of the estimators under several loss functions, such as the least-absolute deviation loss and Huber loss; Wang et al. (2024) proposed an efficient and robust nonparametric regression estimator based on an estimated maximum likelihood estimator; Fan & Gu (2024), Bhattacharya et al. (2024) investigated deep neural network estimators for nonparametric interaction models with diverging input dimensions, focusing on theoretical properties such as minimax optimality, approximation error control, and sparsity-driven regularization. However, quantifying uncertainty of DNN-estimated $\mathbb{E}(y|\mathbf{x})$ within the framework of GNRMs introduces additional challenges. The existing approaches typically assume a nonparametric regression framework, where the distribution of $y - \mathbb{E}(y|\mathbf{x})$ does not depend on \mathbf{x} . This assumption simplifies the theoretical analysis and the derivation of convergence results (Takezawa 2005, Györfi et al. 2006, Fan 2018). In generalized nonparametric regression settings, however, the assumption often fails due to heteroskedasticity, as the distribution of $y - \mathbb{E}(y|\mathbf{x})$ may still depend on \mathbf{x} . Consequently, the standard techniques for bounding error terms and controlling variance no longer apply, and the established results on uniform convergence or concentration inequalities (Bartlett et al. 2020) are invalid. As a result, it is unclear whether these methods are applicable for drawing inference on GNRMs estimated by DNNs.

Our methodological and theoretical contributions are as follows. First, we propose a DNN estimator for subject specific means under GNRMs and develop an inference framework that explicitly accounts for covariate-dependent error distributions induced by heteroskedasticity. We provide new proof techniques, including new criterion loss function, covering-number based concentration and a truncation strategy that handles exponential-tailed, heteroskedastic noise in Section D in the Supplementary Material, improving the results of Schmidt-Hieber (2020). For example, with the general loss induced by GNRMs,

the quadratic-based arguments of [Schmidt-Hieber \(2020\)](#) no longer apply. This necessitates a complete redevelopment of the error bounds, derived instead from the intrinsic curvature and Lipschitz properties in the new GNRM setting. As a result, all estimation error analyses are novel and move beyond Gaussian approximations. Additionally, the presence of exponential tailed, heteroskedastic noise complicates uniform convergence, which we address through a truncation strategy. The truncation scheme partitions the analysis into bounded and unbounded noise regimes, where standard convergence arguments apply in the former and concentration inequalities control the latter, yielding sharper error bounds. These theoretical tools enable us to extend the continuous outcome setting of [Schmidt-Hieber \(2020\)](#), accommodating both continuous and discrete outcomes, and thereby generalizing their framework to the broader GNRM class. Second, to quantify the uncertainty of the estimates, we propose an ensemble subsampling method (ESM) under GNRMs, developed from U-statistics ([Frees 1989](#), [Hoeffding 1992](#), [Lee 2019](#), [Boroskikh 2020](#)) and Hoeffding decomposition ([Hoeffding 1992](#)). Our model-free variance estimator leverages the intrinsic properties of Hoeffding decomposition and accommodates varying population levels as well as individual variance structures. Moreover, we employ incomplete U-statistics to reduce computational cost, enabling the resample size to grow proportionally with the sample size while limiting the total number of resamples. To mitigate the resulting Monte Carlo bias, we refine the variance estimator of [Wager & Athey \(2018\)](#) for bias correction.

In [Section 2](#), we introduce deep neural networks (DNNs), generalized nonparametric regression models (GNRMs) and the ensemble subsampling method (ESM). [Section 3](#) presents theoretical results for the convergence rate of DNNs under the GNRM framework, as well as results for statistical inference using ESM. In [Section 4](#), we perform simulations to evaluate the finite sample performance of our method. [Section 5](#) demonstrates the proposed

methodology using the electronic Intensive Care Unit (eICU) dataset to estimate both the probability and expected number of ICU readmissions based on patient characteristics, with accompanying statistical inference. The eICU dataset is a multi-center resource comprising anonymized clinical data from ICU patients, developed through a collaboration between the Massachusetts Institute of Technology and multiple healthcare partners. The findings may support the identification of high-risk individuals and guide the delivery of personalized, risk-informed care. Section 6 concludes the paper with discussions. Additional experiments and proofs are provided in the Supplementary Material.

Notation. We employ lower letters (e.g. a) for scalars and boldface letters for vectors and matrices (e.g. \mathbf{a} and \mathbf{A}). For any matrix \mathbf{A} , we use $\|\mathbf{A}\|_0$ and $\|\mathbf{A}\|_\infty$ to denote its zero norm and infinity norm, respectively. The set of natural and real numbers are denoted by \mathbb{N} and \mathbb{R} respectively. For any function f mapping a vector to \mathbb{R} , we use $\|f\|_\infty$ to denote its infinity norm. For an integer n , $[n] = \{1, \dots, n\}$. We write $X_1(n) = O(X_2(n))$ or $X_1(n) \lesssim X_2(n)$ or $X_2(n) \gtrsim X_1(n)$ if there exist $C > 0$ and $N_0 > 0$ such that $|X_1(n)| \leq C|X_2(n)|$ when $n > N_0$. We denote $X_2(n) \asymp X_1(n)$ if $X_1(n) = O(X_2(n))$ and $X_2(n) = O(X_1(n))$. We denote $X_1(n) = o(X_2(n))$ if $X_1(n)/X_2(n) \rightarrow 0$, and $X_1(n) \sim X_2(n)$ if $\{X_1(n)/X_2(n)\}$ converges to 1. When clear from context, we add the index p (for probability) such as O_p for a random variable and its realized value to avoid redundancy.

2 The Preamble

2.1 Deep Neural Networks

Let L be a positive integer denoting the **depth** of a neural network, and $\mathbf{p} = (p_0, \dots, p_L, p_{L+1})^\top$ be a vector of positive integers specifying the **width** of each layer, with p_0 corresponding

to the input dimension, p_1, \dots, p_L the dimensions of the L hidden layers, and p_{L+1} the dimension of the output layer. An $(L + 1)$ -layer DNN with layer-width \mathbf{p} is defined as

$$f(\mathbf{x}) = \mathbf{W}_L f_L(\mathbf{x}) + \mathbf{v}_L, \quad (2.1)$$

with the recursive expression $f_l(\mathbf{x}) = \sigma(\mathbf{W}_{l-1} f_{l-1}(\mathbf{x}) + \mathbf{v}_{l-1})$, $f_1(\mathbf{x}) = \sigma(\mathbf{W}_0 \mathbf{x} + \mathbf{v}_0)$. Here, $\mathbf{W}_l \in \mathbb{R}^{p_{l+1} \times p_l}$ and vectors $\mathbf{v}_l \in \mathbb{R}^{p_{l+1}}$ ($l \in [L]$) are the parameters of this DNN, and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function, applied element-wise to vectors. A commonly used activation function is the rectified linear unit (ReLU) function (Nair & Hinton 2010), or $\sigma(x) = \max(x, 0)$. We focus on ReLU due to its piecewise linearity and projection properties, with potential extensions to other activation functions as discussed in Section 6.

Overparameterized DNNs can memorize noise and overfit (Cao et al. 2022), so sparsity is used to control complexity. We focus on the following class of s -sparse and F -bounded networks for our theoretical analysis:

$$\mathcal{F}(L, \mathbf{p}, s, F) = \left\{ f \text{ of form (2.1): } \|\mathbf{W}_l\|_\infty, \|\mathbf{v}_l\|_\infty \leq 1, \sum_{l=1}^L \|\mathbf{W}_l\|_0 + \|\mathbf{v}_l\|_0 \leq s, \|f\|_\infty \leq F \right\}. \quad (2.2)$$

Here, $s \in \mathbb{N}_+$, and $F > 0$ is a constant. $\|\cdot\|_0$ is the number of nonzero entries of matrix or vector, and $\|f\|_\infty$ is the sup-norm of function f .

2.2 Generalized Nonparametric Regression

Consider outcomes from the exponential family, including commonly used distributions such as normal, Bernoulli, Poisson, and binomial. The density (or mass) function for a random variable y in the (single-parameter) exponential family is $p(y|\theta) = h(y) \exp\{\theta y - \psi(\theta)\}$, where θ is the canonical parameter, $0 < h(y) < \infty$ ensures normalization, and $\psi(\theta)$ is convex and smooth, satisfying $\mathbb{E}(y|\theta) = \psi'(\theta)$, where $\psi'(\cdot)$ denotes the derivative of $\psi(\cdot)$ and the inverse function of $\psi'(\cdot)$ is known as the link function (Brown 1986, Fei & Li 2021).

Suppose we observe n independently and identically distributed (iid) sample $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ from (\mathbf{x}, y) , where $\mathbf{x} \in \mathbb{R}^p$ is a covariate vector and $y \in \mathbb{R}$ is the response variable. Given \mathbf{x} , the distribution of y is modeled as

$$p(y|\mathbf{x}) = h(y) \exp \left\{ yf_0(\mathbf{x}) - \psi(f_0(\mathbf{x})) \right\}, \quad (2.3)$$

where $f_0 : \mathbb{R}^p \rightarrow \mathbb{R}$ is a nonparametric function with $f_0(\mathbf{x})$ representing the mean parameter of the response variable y given \mathbf{x} . Model (2.3) is referred to as a Generalized Nonparametric Regression Model (GNRM), where the outcomes follow a distribution from the exponential family. With an independent test point $(\mathbf{x}_{\text{new}}, y_{\text{new}})$ and that y_{new} , given \mathbf{x}_{new} , follows (2.3), we aim to estimate $\mathbb{E}(y_{\text{new}}|\mathbf{x}_{\text{new}})$ which is $\psi'(f_0(\mathbf{x}_{\text{new}}))$. Consequently, it is of interest to estimate f_0 . By the universal approximation theorem (Hornik et al. 1989), DNNs are one of suitable candidates for estimating f_0 . We consider an optimal DNN estimator, \hat{f}_n^{opt} , which minimizes the negative log likelihood function (2.3) over the class of DNNs defined in (2.2), that is,

$$\hat{f}_n^{\text{opt}} = \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n \{-y_i f(\mathbf{x}_i) + \psi(f(\mathbf{x}_i))\}. \quad (2.4)$$

The estimator (2.4) is applicable to commonly used models, including these listed below.

Nonparametric Gaussian Regression: Given \mathbf{x} , we assume that the response variable follows a Gaussian distribution with mean $f_0(\mathbf{x})$ and variance σ^2 . In this setting, σ^2 does not affect the estimation of the mean function $f_0(\mathbf{x})$, which can be seen by examining the likelihood and optimization. The likelihood function for a nonparametric Gaussian regression model is given by $p(y|\mathbf{x}) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(y-f_0(\mathbf{x}))^2}{2\sigma^2}}$. An application of (2.4) to this setting yields:

$$\begin{aligned} \hat{f}_n^{\text{opt}} &= \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n \{-y_i f(\mathbf{x}_i) + \frac{1}{2} f^2(\mathbf{x}_i)\} \\ &= \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n \{y_i - f(\mathbf{x}_i)\}^2, \end{aligned}$$

corresponding to mean squared error minimization as in [Schmidt-Hieber \(2020\)](#).

Nonparametric Logistic Regression: To model binary outcomes $y \in \{0, 1\}$, we set $h(y) = 1$, $\psi(\theta) = \log(1 + e^\theta)$ in (2.3). The estimator (2.4) in this setting can be written as:

$$\widehat{f}_n^{\text{opt}} = \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n (-y_i f(\mathbf{x}_i) + \log[1 + \exp\{f(\mathbf{x}_i)\}]).$$

Nonparametric Poisson Regression: Considering $y \in \mathbb{N}$ (non-negative integers), we define $h(y) = \frac{1}{y!}$ and $\psi(\theta) = e^\theta$ in (2.3). The estimator (2.4) has the following form:

$$\widehat{f}_n^{\text{opt}} = \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n [\exp\{f(\mathbf{x}_i)\} - y_i f(\mathbf{x}_i)].$$

Nonparametric Binomial Regression: Given covariates \mathbf{x} , the response variable y follows a Binomial distribution with a fixed number of trials n_{trial} and the success probability $p(\mathbf{x}) = \frac{1}{1 + e^{-f_0(\mathbf{x})}}$. This corresponds to $h(y) = \binom{n_{\text{trial}}}{y}$ and $\psi(\theta) = n_{\text{trial}} \log(1 + e^\theta)$. In this setting, the estimator (2.4) is specified as

$$\widehat{f}_n^{\text{opt}} = \operatorname{argmin}_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \frac{1}{n} \sum_{i=1}^n \left[-y_i f(\mathbf{x}_i) + n_{\text{trial}} \log\{1 + e^{f(\mathbf{x}_i)}\} \right].$$

2.3 Empirical Risk Minimization and Function Smoothness

Recall that the pointwise Kullback–Leibler divergence between two distributions of the form (2.3), one specified by f and another by the true f_0 , conditional on \mathbf{x} , i.e.,

$$\mathbb{E} \left[-yf(\mathbf{x}) + \psi(f(\mathbf{x})) + yf_0(\mathbf{x}) - \psi(f_0(\mathbf{x})) \mid \mathbf{x} \right] = -\psi'(f_0(\mathbf{x}))(f(\mathbf{x}) - f_0(\mathbf{x})) + \psi(f(\mathbf{x})) - \psi(f_0(\mathbf{x})),$$

where we use $\mathbb{E}[y|\mathbf{x}] = \psi'(f_0(\mathbf{x}))$. This equals the Bregman divergence between $f(\mathbf{x})$ and $f_0(\mathbf{x})$, which is nonnegative and vanishes if and only if $f(\mathbf{x}) = f_0(\mathbf{x})$. We then quantify the estimation error of \widehat{f} , any estimator of f_0 , by evaluating the pointwise estimation error at an independent test point \mathbf{X} : $\ell(\mathbf{X}; \widehat{f}, f_0) = -\psi'(f_0(\mathbf{X}))(\widehat{f}(\mathbf{X}) - f_0(\mathbf{X})) + \psi(\widehat{f}(\mathbf{X})) -$

$\psi(f_0(\mathbf{X}))$. We further define the average estimation error, our main theoretical target, as

$$R_n(\hat{f}, f_0) := \mathbb{E}[\ell(\mathbf{X}; \hat{f}, f_0)], \quad (2.5)$$

where the expectation is taken with respect to all involved random variables.

Computing the exact minimizer \hat{f}_n^{opt} in (2.4) is challenging due to the nonconvex loss (Fan & Gu 2024). Instead, we obtain an approximate minimizer \hat{f}_n using optimization algorithms such as gradient descent. To quantify the optimization error, we introduce

$$\Delta_n(\hat{f}_n, \hat{f}_n^{\text{opt}}) = \mathbb{E} \left\{ \frac{1}{n} \sum_{i=1}^n (-y_i \hat{f}_n(\mathbf{x}_i) + \psi(\hat{f}_n(\mathbf{x}_i))) - \frac{1}{n} \sum_{i=1}^n (-y_i \hat{f}_n^{\text{opt}}(\mathbf{x}_i) + \psi(\hat{f}_n^{\text{opt}}(\mathbf{x}_i))) \right\} \quad (2.6)$$

that measures the difference between the expected empirical risk of \hat{f}_n and \hat{f}_n^{opt} . For notational ease, we write $\Delta_n(\hat{f}_n) := \Delta_n(\hat{f}_n, \hat{f}_n^{\text{opt}})$ when the reference to \hat{f}_n^{opt} is clear. We will show that a small empirical deviation $\Delta_n(\hat{f}_n)$, combined with controlled model complexity of the neural network class, implies a bound on the average excess risk $R_n(\hat{f}_n, f_0)$, forming the foundation for the inference theory developed in subsequent sections.

For desirable properties of \hat{f}_n , we assume f_0 belongs to a Hölder class (Schmidt-Hieber 2020, Fan & Gu 2024) with parameters $\gamma, K > 0$, and domain $\mathbb{D} \subset \mathbb{R}^r$:

$$\mathcal{G}_r^\gamma(\mathbb{D}, K) = \left\{ g : \mathbb{D} \rightarrow \mathbb{R} : \sum_{\beta: \|\beta\|_1 < \gamma} \|\partial^\beta g\|_\infty + \sum_{\beta: \|\beta\|_1 = \lfloor \gamma \rfloor} \sup_{\mathbf{x}, \mathbf{y} \in \mathbb{D}, \mathbf{x} \neq \mathbf{y}} \frac{|\partial^\beta g(\mathbf{x}) - \partial^\beta g(\mathbf{y})|}{\|\mathbf{x} - \mathbf{y}\|_\infty^{\gamma - \lfloor \gamma \rfloor}} \leq K \right\},$$

where $\lfloor \gamma \rfloor$ is the largest integer strictly smaller than γ , and $\partial^\beta = \partial^{\beta_1} \dots \partial^{\beta_r}$ with $\beta = (\beta_1, \dots, \beta_r)^\top \in \mathbb{N}^r$. We further assume that f_0 is a $(q+1)$ -composition ($q \in \mathbb{N}$) of several Hölder functions. That is, for some vectors $\mathbf{d} = (d_0, \dots, d_{q+1}) \in \mathbb{N}_+^{q+2}$, $\mathbf{t} = (t_0, \dots, t_q) \in \mathbb{N}_+^{q+1}$ and $\mathbf{m} = (m_0, \dots, m_q) \in \mathbb{R}_+^{q+1}$, $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \mathbf{m}, K)$, where

$$\mathcal{G}(q, \mathbf{d}, \mathbf{t}, \mathbf{m}, K) := \left\{ f = g_q \circ \dots \circ g_0 : g_i = (g_{ij})_j : [a_i, b_i]^{d_i} \rightarrow [a_{i+1}, b_{i+1}]^{d_{i+1}}, \right. \\ \left. g_{ij} \in \mathcal{G}_{t_i}^{m_i}([a_i, b_i]^{t_i}) \text{ with } |a_i|, |b_i| \leq K \right\}.$$

The motivation for the assumptions on f_0 is that the Hölder class provides a flexible framework for quantifying smoothness, capturing both integer and fractional orders, and serves as a natural benchmark for assessing neural network approximation (Yarotsky 2017), while the compositional structure mirrors the architectural design of deep neural networks with each layer applying a transformation to the output of the previous one (Yarotsky 2017). We also note that \mathbf{d} and \mathbf{t} characterize the dimensions of the function, while \mathbf{m} represents the measure of smoothness. For instance, if

$$f_0(\mathbf{x}) = g_{21} \left(g_{11} \left(g_{01}(x_1, x_2, x_3, x_4), g_{02}((x_5, x_6, x_7, x_8)) \right), \right. \\ \left. g_{12} \left(g_{03}(x_9, x_{10}, x_{11}, x_{12}), g_{04}((x_{13}, x_{14}, x_{15}, x_{16})) \right), \right. \\ \left. g_{13} \left(g_{05}(x_{17}, x_{18}, x_{19}, x_{20}), g_{06}((x_{21}, x_{22}, x_{23}, x_{24})) \right) \right), \quad \mathbf{x} \in [0, 1]^{24},$$

and g_{ij} are twice continuously differentiable, then smoothness $\mathbf{m} = (2, 2, 2)$, dimensions $\mathbf{d} = (24, 6, 3, 1)^\top$ and $\mathbf{t} = (4, 2, 3)^\top$. We further define $m_j^* = m_j \prod_{l=i+1}^q (m_l \wedge 1)$ and $\phi_n = \max_{j=0, \dots, q} n^{-2m_j^*/(2m_j^* + t_j)}$.

2.4 Ensemble Subsampling Method (ESM)

Consider a scenario where \mathbf{x} is drawn from an unknown probability measure $\mathbb{P}_{\mathbf{x}}$ supported on \mathcal{X} , and for an input $\mathbf{x}_{\text{new}} = \mathbf{x}_* \in \mathcal{X}$, the unobserved outcome y_{new} follows (2.3). Our goal is to estimate $\mathbb{E}(y_{\text{new}} | \mathbf{x}_{\text{new}} = \mathbf{x}_*) = \psi'(f_0(\mathbf{x}_*))$ and its uncertainty. For this, we introduce the Ensemble Subsampling Method (ESM), based on subsampling techniques (Wager & Athey 2018, Fei & Li 2024), to construct ensemble estimators and confidence intervals. ESM consists of the following components, which are summarized in Figure 1 as well.

- (Subsampling) Let $\mathcal{I} = \{1, \dots, n\}$ denote the index set of the training dataset \mathcal{D}_n . We construct subsets of size $r (< n)$, yielding $B^* = \binom{n}{r}$ unique combinations. Denoting the b -th subset as $\mathcal{I}^b = \{i_1, \dots, i_r\}$, where $i_1 < \dots < i_r$ and $b \in [1 : B^*]$.

- (Estimator Construction) For each $1 \leq b \leq B^*$, we aim to minimize the objective function in (2.4) within $\mathcal{F}(L, \mathbf{p}, s, F)$ by using a gradient based optimization algorithm such as Stochastic Gradient Descent (SGD) on the observations indexed by \mathcal{I}^b . Let \hat{f}^b denote the resulting approximate minimizer satisfying

$$\Delta_r(\hat{f}^b, \hat{f}^{b,\text{opt}}) \leq \Delta_b^{\text{opt}}, \quad (2.7)$$

where Δ_b^{opt} quantifies the optimization error (Fan & Gu 2024), and $\hat{f}^{b,\text{opt}}$ and $\Delta_r(\hat{f}^b, \hat{f}^{b,\text{opt}})$ are defined respectively in (2.4) and (2.6) on the subsample indexed by \mathcal{I}^b .

- (Ensemble Estimation) Using the estimators $\{\hat{f}^b : b = 1, \dots, B^*\}$, we compute the ensemble estimation for any $\mathbf{x}_* \in \mathcal{X}$ as:

$$\hat{f}^{B^*}(\mathbf{x}_*) = \frac{1}{B^*} \sum_{b=1}^{B^*} \hat{f}^b(\mathbf{x}_*),$$

and estimate $\mathbb{E}(y_{\text{new}} | \mathbf{x}_{\text{new}} = \mathbf{x}_*)$ as $\psi'(\hat{f}^{B^*}(\mathbf{x}_*))$.

- (Confidence Interval Construction) To quantify uncertainty, we estimate the variance, denoted $\hat{\sigma}_*$, using the infinitesimal jackknife method. The resulting confidence interval (CI) is defined as:

$$\text{CI}(\mathbf{x}_*) = [\psi'(\hat{f}^{B^*}(\mathbf{x}_*) - c_\alpha \hat{\sigma}_*), \psi'(\hat{f}^{B^*}(\mathbf{x}_*) + c_\alpha \hat{\sigma}_*)]. \quad (2.8)$$

where c_α is a constant controlling the $1 - \alpha$ confidence level.

The distribution of $\hat{f}^b(\mathbf{x}_*)$ for a fixed \mathbf{x}_* is intractable. However, since $\hat{f}^b(\mathbf{x}_*)$ is permutation symmetric with respect to \mathcal{I}^b , the ensemble estimator $\hat{f}^{B^*}(\mathbf{x}_*)$ forms a generalized U-statistic, which we later show to be asymptotically normal by leveraging U-statistics theory. In practice, we allow r to grow with n , but evaluating all $B^* = \binom{n}{r}$ models becomes computationally infeasible. To overcome this, we use a stochastic approximation by randomly drawing B size- r subsamples from \mathcal{D}_n . This involves independently sampling indices

ESM Pipeline

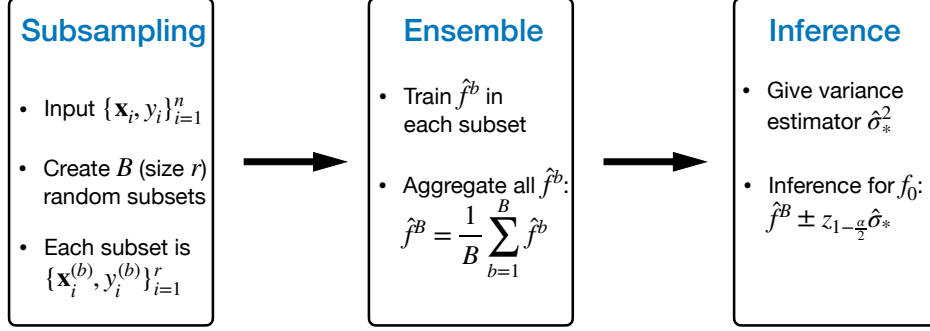


Figure 1: Overview of the Ensemble Subsampling Method (ESM).

b_1, \dots, b_B from 1 to $\binom{n}{r}$, where each subsample is indexed by \mathcal{I}^{b_j} , and computing

$$\hat{f}^B(\mathbf{x}_*) = \frac{1}{B} \sum_{j=1}^B \hat{f}^{b_j}(\mathbf{x}_*). \quad (2.9)$$

As (2.9) provides the estimator for f_0 , the variance $\hat{\sigma}_*$ in (2.8) in the Confidence Interval Construction procedure is given as

$$\hat{\sigma}_*^2 = \underbrace{\frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \hat{V}_i^2}_{\text{original } \hat{\sigma}^2} - \underbrace{\frac{n(n-1)}{(n-r)^2} \cdot \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B (Z_{b_j i} - \hat{V}_i)^2}_{\text{bias correction}}. \quad (2.10)$$

Here, $Z_{b_j i}$ and \hat{V}_i are defined as

$$Z_{b_j i} = (J_{b_j i} - J_{\cdot i})(\hat{f}^{b_j}(\mathbf{x}_*) - \hat{f}^B(\mathbf{x}_*)), \quad \hat{V}_i = \frac{\sum_{j=1}^B Z_{b_j i}}{B},$$

where $J_{b_j i} = I(i \in \mathcal{I}^{b_j})$ and $J_{\cdot i} = \frac{1}{B} \sum_{j=1}^B J_{b_j i}$. Extending Peng et al. (2024), (2.10) jointly corrects Monte Carlo bias and finite-sample effects using the factor $n(n-1)/(n-r)^2$, enabling valid inference with $B \gtrsim n$. Its construction relies on the dominance of the first order term in the Hoeffding decomposition, and the bias correction term adequately captures its resampling variability, regardless of the specific scaling between r and n .

3 Theoretical Results

Let $\mathbf{z}_i = (y_i, \mathbf{x}_i)$ represent an independent observation of sample points. For studying the asymptotic properties of \widehat{f}^B in (2.9), we introduce

$$\xi_{1,r}(\mathbf{x}) = \text{Cov}(\widehat{f}(\mathbf{x}; \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r), \widehat{f}(\mathbf{x}; \mathbf{z}_1, \mathbf{z}'_2, \dots, \mathbf{z}'_r)). \quad (3.1)$$

Here $\mathbf{x} \in \mathcal{X}$ is in the support of probability measure $\mathbb{P}_{\mathbf{x}}$, $\widehat{f}(\mathbf{x}; \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r)$ means that we obtain \widehat{f} from the subsample $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r$ and then apply it to the point \mathbf{x} , and $\xi_{1,r}$ quantifies the covariance between estimates based on two overlapping subsets of size r that share a common point \mathbf{z}_1 but differ in the remaining elements. $\xi_{1,r}$ represents a second-order component in the Hoeffding decomposition of a U-statistic (Hoeffding 1992), capturing the pairwise dependence structure within the subsamples used for constructing the estimator. The terms \mathbf{z}'_i and \mathbf{z}_i are independently generated from the same data generation process. We impose some regularity assumptions for our theoretical guarantees.

Assumption 3.1. *Suppose that (2.3) holds and the true function $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \mathbf{m}, K)$. The support of y , denoted by $\text{supp}(h(y))$, is fixed and does not depend on \mathbf{x} . Moreover, the sufficient statistic $T(y) = y$ is not degenerate and can take at least two distinct values in $\text{supp}(h)$. For any \mathbf{x} , the conditional distribution $p(y|\mathbf{x})$ is sub-exponential. Specifically, with bounded f_0 there exists a universal constant $\kappa > 0$ such that for all $t > 0$,*

$$\mathbb{P}(|y - \mathbb{E}(y|\mathbf{x})| \geq t|\mathbf{x}) \leq 2 \exp\left(-\frac{t}{\kappa}\right).$$

Assumption 3.2. *The network class $\mathcal{F}(L, \mathbf{p}, s, F)$ with L , \mathbf{p} and s satisfies i) $F \geq \max(K, 1)$, ii) $\sum_{i=0}^q \log_2(4t_i \vee 4m_i) \log_2 n \leq L \lesssim \log^\alpha n$ for some $\alpha > 1$, iii) $n\phi_n \lesssim \min_{i=1, \dots, L} p_i \leq \max_{i=1, \dots, L} p_i \lesssim n$, iv) $s \lesssim n\phi_n \log n$. Here, $\phi_n = \max_{j=0, \dots, q} n^{-2m_j^*/(2m_j^*+t_j)}$, where $m_j^* = m_j \prod_{l=i+1}^q (m_l \wedge 1)$. The values t_i, m_i and p_i are the $(i+1)$ -th element in the vectors \mathbf{t} , \mathbf{m} and \mathbf{p} , respectively.*

Assumption 3.3. For each subsample with index \mathcal{I}^b and size r , there exists a constant $C > 0$ such that the optimization error Δ_b^{opt} defined in (2.7) satisfies $\Delta_b^{\text{opt}} \leq C\phi_r L \log^2 r$.

Assumption 3.4. The covariance $\xi_{1,r}(\mathbf{x})$ defined in (3.1) satisfies $\inf_{\mathbf{x} \in \mathcal{X}} \xi_{1,r}(\mathbf{x}) \geq Cr^{-(1+\varepsilon)}$ for some $\varepsilon > 0$. The subsample size $r = n^\gamma$ satisfies $1/(1 + \min_{j=0,\dots,q}(2m_j^*)/(2m_j^* + t_j)) < \gamma < 1$, and the number of subsampling times B is large such that $B \gtrsim n$.

Assumption 3.1 ensures that $p(y|\mathbf{x})$ not degenerate, while the concentration of y does not deviate too widely. Additionally, that $T(y) = y$ takes at least two distinct values in $\text{supp}(h)$ guarantees a non-zero variance of y , ensuring the log-partition function $\psi(f_0(\mathbf{x}))$ exhibits local convexity. The boundedness of f_0 guarantees the existence of κ . This condition, weaker than the sub-Gaussian tail bound assumed in Fan et al. (2024), is satisfied by most common exponential family distributions, including Beta, Bernoulli, Poisson, and Gaussian. Assumption 3.2 is a technical assumption similar to Schmidt-Hieber (2020), which specifies the requirements on the network parameters relative to the sample size n and smoothness parameters ϕ_n , t_i and m_i , including the number of layers L and the sparsity s . Assumption 3.3, related to the training dynamics of DNNs, stipulates that DNNs return an estimator sufficiently close to the global minimum. We note that the rate in Assumption 3.3 pertains to the prediction error in Theorem 3.5. Our intent is not to assert an explicit convergence rate for SGD in practice, but to impose a sufficient condition that separates optimization error from the inferential analysis. This perspective follows standard practice in the statistical literature (Chernozhukov et al. 2018, Schmidt-Hieber 2020, Fan & Gu 2024), where such assumptions serve as technical devices to ensure valid inference rather than as primary objects of study. Assumption 3.4 on $\xi_{1,r}(\mathbf{x})$ is not meant to serve as a precise characterization of neural networks. Rather, it ensures that the first-order Hoeffding projection remains non-degenerate and dominates the higher-order components,

as in the random forest setting (Wager & Athey 2018). We choose the subsample size r much smaller than n to ensure first-order dominance but large enough to limit bias, and let the number of subsamples B scale with n to control computation.

Theorem 3.5. *Suppose that $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \mathbf{m}, K)$ and $\widehat{f}_n \in \mathcal{F}(L, \mathbf{p}, s, F)$. If Assumptions 3.1-3.2 hold, then it holds that*

$$\frac{1}{2}\Delta_n(\widehat{f}_n) - c'\phi_n L \log^2 n \leq R_n(\widehat{f}_n, f_0) \leq 2\Delta_n(\widehat{f}_n) + C'\phi_n L \log^2 n.$$

for some constants $c', C' > 0$. Here, $\Delta_n(\widehat{f}_n)$ is defined in (2.6).

Theorem 3.5 establishes a bound on the out-of-sample estimation error. Unlike previous studies (Schmidt-Hieber 2020, Fan & Gu 2024), which primarily focus on bounding the in-sample mean squared error, i.e., $\frac{1}{n} \sum_{i=1}^n (y_i - f(\mathbf{x}_i))^2$, to derive bounds for the estimation error, our approach applies to a wide class of training losses (2.6) under GNRMs. The heteroskedasticity inherent in GNRMs necessitates a more refined analysis, and a truncation technique is introduced to address this challenge by separating the noise into different regimes. Specifically, this result enables the analysis of non-Gaussian outcomes with heteroskedasticity depending on covariates, overcoming key limitations of prior methods (Schmidt-Hieber 2020, Fan & Gu 2024) that rely on the independence of $y_i - \mathbb{E}(y_i|\mathbf{x}_i)$ and $f_0(\mathbf{x}_i)$ to get estimation error bounds.

Proposition 3.6. *We suppose that the conditions of Theorem 3.5 hold, the neural network fitted on subsamples belongs to $\mathcal{F}(L, \mathbf{p}, s, F)$, Assumptions 3.2-3.4 hold for each subsample with size r , and sparsity $s \asymp r^{1+\delta}\phi_r \log r$ for some constant $\delta > 0$. Then for a fixed \mathbf{x}_* , if the pointwise bias is dominated by the best approximation bias, i.e., there exists constant $C(\mathbf{x}_*)$ only depends on \mathbf{x}_* such that $|\mathbb{E}\widehat{f}^b(\mathbf{x}_*) - f_0(\mathbf{x}_*)| \lesssim C(\mathbf{x}_*) \inf_{f \in \mathcal{F}} |f(\mathbf{x}_*) - f_0(\mathbf{x}_*)|$, it holds that $\sqrt{\frac{n}{r^2 \inf_{\mathbf{x} \in \mathcal{X}} \xi_{1,r}(\mathbf{x})}} \cdot |\mathbb{E}\widehat{f}^B(\mathbf{x}_*) - f_0(\mathbf{x}_*)| \rightarrow 0$.*

Proposition 3.6 implies asymptotic unbiasedness, i.e., $\mathbb{E}\{\widehat{f}^B(\mathbf{x}_*)\} - f_0(\mathbf{x}_*) = o\left(\sqrt{\frac{r^2 \xi_{1,r}(\mathbf{x}_*)}{n}}\right)$ for a fixed \mathbf{x}_* . Thus, the bias of \widehat{f}^B is asymptotically negligible, and enables us to focus on $\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\{\widehat{f}^B(\mathbf{x}_*)\}$ and establish the following results.

Theorem 3.7. *Suppose the data \mathcal{D}_n are generated from Model (2.3) with an unknown $f_0 \in \mathcal{G}(q, \mathbf{d}, \mathbf{t}, \mathbf{m}, K)$, the used neural networks belong to $\mathcal{F}(L, \mathbf{p}, s, F)$, and Assumptions 3.1–3.4 hold (Assumption 3.2 holds for each subsample with size r). Then with the results of Proposition 3.6, there exists a positive sequence, $\delta_1, \delta_2, \dots$, with $\delta_n \rightarrow 0$ and a corresponding set $\mathcal{A}_{\delta_n} \subseteq \mathcal{X}$ with $\mathbb{P}_{\mathbf{x}}(\mathcal{A}_{\delta_n}) \geq 1 - \delta_n$, such that for any fixed $\mathbf{x}_* \in \mathcal{A}_{\delta_n}$, it holds that*

$$\sqrt{n} \cdot \frac{\widehat{f}^B(\mathbf{x}_*) - f_0(\mathbf{x}_*)}{\sqrt{r^2 \xi_{1,r}(\mathbf{x}_*)}} \xrightarrow{d} \mathcal{N}(0, 1).$$

Theorem 3.7 establishes the asymptotic normality of the ensembled estimator \widehat{f}^B for any fixed $\mathbf{x}_* \in \mathcal{A}_{\delta_n}$, i.e., the inference result holds with high probability for almost all $\mathbf{x} \in \mathcal{X}$. Here, Theorem 3.7 holds by the fact that the first-order projection of $\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\{\widehat{f}^B(\mathbf{x}_*)\}$ dominates the other terms. Moreover, as the analytic form of $\xi_{1,r}(\mathbf{x}_*)$ is unavailable, the following theorem provides a consistent estimator of it.

Theorem 3.8. *Under the same setting as in Theorem 3.7 and with $\widehat{\sigma}_*^2$ defined in (2.10), it holds that for any fixed $\mathbf{x}_* \in \mathcal{X}$,*

$$\sqrt{\frac{r^2 \xi_{1,r}(\mathbf{x}_*)}{n \widehat{\sigma}_*^2}} \xrightarrow{p} 1.$$

Compared to Wager & Athey (2018) and Fei & Li (2024), since we allow the number of subsamples B to grow at the same rate as the sample size n , the Monte Carlo bias in the variance estimation becomes non-negligible. As pointed out by Wager et al. (2014), this bias is typically of order n/B , and does not vanish when $B \asymp n$. Different from the correction proposed in Wager et al. (2014), which assumes near asymptotic independence

between the inclusion indicators J_{b_i} and the estimates \hat{f}^b , we propose a bias-correction term that directly accounts for the dependency between subsample structure and estimates. Specifically, we utilize the covariance form of the variance estimator and subtract a U-statistic-based empirical correction that estimates the within-subject variability due to repeated subsampling. This yields a debiased estimator that is asymptotically unbiased even when $B \asymp n$. For the confidence interval in (2.8), Theorem 3.8 and the Slutsky theorem ensure that, as $n \rightarrow \infty$, it holds that

$$\mathbb{P}\left(\mathbb{E}(y_{\text{new}}|\mathbf{x}_*) \in \left[\hat{L}(\mathbf{x}_*), \hat{U}(\mathbf{x}_*)\right]\right) \rightarrow 1 - \alpha,$$

where $\hat{L}(\mathbf{x}_*) = \psi'\left(\hat{f}^B(\mathbf{x}_*) - z_{1-\alpha/2}\hat{\sigma}_*\right)$, $\hat{U}(\mathbf{x}_*) = \psi'\left(\hat{f}^B(\mathbf{x}_*) + z_{1-\alpha/2}\hat{\sigma}_*\right)$, and $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -th quantile of the standard normal distribution. In practice, we recommend using this back-transformed interval, as it preserves the parameter's natural range.

Remark. Our inference framework relies on a well-trained network; when training is insufficient or overly regularized, inference may become unreliable. For analytical tractability, we adopt an idealized s -sparse, F -bounded network class (Schmidt-Hieber 2020), while in practice, regularization techniques such as weight decay and dropout promote approximate sparsity, keeping the trained model close to this class. The optimization gap $\Delta_n(\hat{f}_n)$ in Theorem 3.5 quantifies the discrepancy between empirical and population risks, filters out degenerate training cases, and accounts for residual optimization and approximation errors. Our theory demonstrates that as long as $\Delta_n(\hat{f}_n)$ remains small, inference validity is preserved, consistent with Chernozhukov et al. (2018), Fan & Gu (2024).

4 Numerical Experiments

We conduct simulations under logistic and Poisson generalized nonparametric regression models to evaluate the performance of the ensemble subsampling method (ESM) in point

estimation, variance estimation ($\widehat{\sigma}_*$), coverage, and interval length. For comparison, we include several alternative methods capable of producing confidence intervals, including HulC (Kuchibhotla et al. 2024), Bayesian Neural Networks (Sun et al. 2021), Ensemble Methods (Lakshminarayanan et al. 2017), and the naive bootstrap. Additional experiments are presented in the Supplementary Material, namely, Section A.1 (binomial regression), Section A.2 (comparison with random forests (Breiman 2001)), Section A.3 (kernel regression), Section A.4 (varying covariate dimensions and nonlinearities), and Sections A.5–A.9 (sensitivity analyses for optimization parameters, subsample size r , number of subsamples B , network depth L , and dropout rate).

We present our basic settings. Define $g(\mathbf{x}) = x_1 + 0.25x_2^2 + 0.1 \arctan(0.5x_3 - 0.3)$. Here, x_j is the j -th element in the vector \mathbf{x} . We generate \mathbf{x}_i from $\mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ with $p = 10$, i.e., the first 3 covariates are signals and the rest are noise variables. For each $i = 1, \dots, n$ and given \mathbf{x}_i , we independently simulate y_i under the following models.

(logistic model) $\mathbb{P}(y_i = 1 | \mathbf{x}_i) = \frac{1}{1 + \exp\{-g(\mathbf{x}_i)\}}$, with $f_0(\mathbf{x}) = g(\mathbf{x})$ and $\psi'(x) = (1 + \exp(-x))^{-1}$.

(Poisson model) With $k \in \{0, 1, 2, \dots\}$, $\mathbb{P}(y_i = k | \mathbf{x}_i) = \frac{e^{-\lambda(\mathbf{x}_i)} (\lambda(\mathbf{x}_i))^k}{k!}$, where we set $\lambda(\mathbf{x}_i) = \log(1 + e^{g(\mathbf{x}_i)})$ to ensure that $\lambda(\mathbf{x}_i) > 0$. In this model, we have $f_0(\mathbf{x}_i) = \log\{\lambda(\mathbf{x}_i)\}$ and $\psi'(x) = \exp(x)$.


For evaluation, we generate a total of 80 independent test points $\mathbf{x}_{\text{test},i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{10})$, $i = 1, \dots, 80$, which remain fixed throughout the following experiment.

We then apply (2.9) and (2.10) to obtain the point estimate $\widehat{f}^B(\mathbf{x}_{\text{test},i})$ and its associated variance estimate $\widehat{\sigma}_*(\mathbf{x}_{\text{test},i})$. To assess the accuracy of the variance estimates and the empirical coverage of the confidence intervals (2.8), we repeat the entire procedure 300 times, each time independently regenerating the training data. This yields replicated estimates $\widehat{f}_s^B(\mathbf{x}_{\text{test},i})$ and $\widehat{\sigma}_{*,s}(\mathbf{x}_{\text{test},i})$ for $s = 1, \dots, 300$. Final evaluation is based on the average performance across the fixed test set, using the metrics summarized in Table 1.

Table 1: Simulation results for different regressions under varying sample sizes n and subsampling ratios r using the following metrics: absolute Bias_f and MAE_f are, respectively, the average and mean absolute bias between \widehat{f}_s^B and f_0 ; $\text{Bias}_{\psi'}$ and $\text{MAE}_{\psi'}$ are, respectively, the average and mean absolute bias in the transformed form $\psi'(\widehat{f}_s^B)$ and $\psi'(f_0)$; EmpSD denotes the empirical standard deviation of $\widehat{f}_s^B(\mathbf{x}_{\text{test},i})$ computed across the 300 repetitions and then averaged over the 80 test points; SE and SE_c denote the average standard errors without and with bias correction, respectively, averaged over the 80 test points; CP is the average, across test points, of the coverage probability that the 95% confidence interval (2.9) contains the truth over 300 repetitions, and AIL is the mean length of these intervals.

Numbers in brackets indicate the standard deviation of each metric across test points.

	Bias_f	MAE_f	$\text{Bias}_{\psi'}$	$\text{MAE}_{\psi'}$	EmpSD	SE	SE_c	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.11(0.72)	0.57(0.44)	0.02(0.14)	0.11(0.08)	0.65	0.78	0.66	91.7%	0.46(0.13)
$r = n^{0.85}$	0.06(0.59)	0.47(0.36)	0.01(0.12)	0.10(0.07)	0.55	0.68	0.58	94.0%	0.43(0.11)
$r = n^{0.9}$	0.07(0.50)	0.40(0.31)	0.01(0.10)	0.08(0.06)	0.48	0.59	0.50	94.6%	0.39(0.09)
$n = 700, r = n^{0.8}$	0.08(0.41)	0.33(0.26)	0.02(0.09)	0.07(0.05)	0.39	0.51	0.39	93.6%	0.31(0.08)
$r = n^{0.85}$	0.05(0.36)	0.28(0.22)	0.01(0.08)	0.06(0.05)	0.34	0.44	0.34	93.9%	0.28(0.06)
$r = n^{0.9}$	0.06(0.33)	0.26(0.21)	0.01(0.07)	0.06(0.05)	0.31	0.40	0.31	92.9%	0.26(0.06)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.32(0.38)	0.38(0.31)	-0.16(0.23)	0.22(0.18)	0.35	0.43	0.36	88.6%	0.86(0.45)
$r = n^{0.85}$	-0.23(0.36)	0.33(0.27)	-0.12(0.25)	0.21(0.18)	0.34	0.43	0.36	91.9%	0.95(0.52)
$r = n^{0.9}$	-0.14(0.34)	0.29(0.24)	-0.06(0.25)	0.19(0.18)	0.33	0.42	0.35	94.3%	1.02(0.59)
$n = 700, r = n^{0.8}$	-0.17(0.26)	0.25(0.20)	-0.10(0.19)	0.16(0.14)	0.25	0.34	0.26	90.7%	0.67(0.34)
$r = n^{0.85}$	-0.10(0.26)	0.21(0.17)	-0.06(0.19)	0.15(0.14)	0.24	0.33	0.25	93.2%	0.71(0.38)
$r = n^{0.9}$	-0.04(0.24)	0.19(0.15)	-0.02(0.19)	0.14(0.13)	0.23	0.32	0.24	94.5%	0.72(0.38)



(a) Point Estimates and Variations

(b) Estimated Standard Errors

Figure 2: Estimation and inference in simulation samples: Logistic Model with $n = 700$, $r = n^{0.9}$ and $B = 1400$. Figure 2a shows the average estimated $\mathbb{E}(y|\mathbf{x})$ with variability across 300 repetitions (gray band) over test points. Figure 2b compares corrected and uncorrected standard errors and their variability (gray and blue bands, respectively) to the empirical standard deviations of estimates across all test samples.

In the experiment, we randomly select sub-datasets of size r to train neural networks, varying r from $n^{0.8}$ to $n^{0.9}$, with $n = 400$ and 700 . These choices of r are based on our experiments as they provide enough data for neural network convergence while remaining asymptotically smaller than the total sample size n . To ensure computational feasibility, we set $B = 1400$ and use a three-layer DNN with architecture $\mathbf{p} = (p, 128, 64, 1)$, a learning rate of 0.1, and ReLU activation. We apply early stopping (500 epochs), a 10% dropout rate, ℓ_2 weight decay with rate 0.02, and output truncation ($F = 3$) to effectively control the sparsity s and boundedness F in the function class defined in (2.2), aligning the implemented network with theoretical assumptions.

Table 1 shows that increasing the sample size n reduces biases with decreased mean absolute errors (MAE) for both f and its transformed counterpart ψ' . Figure 2a displays the average estimated $\mathbb{E}(y|\mathbf{x})$ and its variation across 300 repetitions for various test points, in the Bernoulli case with $n = 700$ and $r = n^{0.9}$. The estimated means closely align with the true values. To further assess the quality of uncertainty quantification, we also examine the empirical standard deviation across repetitions. The corrected standard errors SE_c align much more closely with the empirical standard deviations (EmpSD) than the uncorrected standard errors (SE), confirming the effectiveness of the correction term introduced in (2.10). This is further illustrated in Figure 2b, where the variance estimator with bias correction closely matches the empirical standard deviation, while the uncorrected version overestimates the uncertainty. As a result, our method achieves coverage probabilities (CP) close to the nominal level.

We present the comparison results in Table 2. For ESM, we select our results with $\gamma = 0.9$. For HulC method proposed by Kuchibhotla et al. (2024), although the coverage probability is close to the nominal level of 95%, its average interval length is very wide. For Bayesian neural networks (Sun et al. 2021), we observe that the performance is highly sensitive to the tuning of hyperparameters. Due to the difficulty of selecting appropriate priors and optimization settings, the coverage probabilities and average interval lengths in our experiments are worse than those of the proposed ESM, with noticeably wider intervals and less stable coverage. For the ensemble methods (Lakshminarayanan et al. 2017) and the naive bootstrap approach, the intervals are wider than those of ESM. As the sample size n increases, and the coverage probability tends to exceed the nominal 95%. Overall, ESM achieves competitive coverage performance with narrowest intervals, although it may not always provide coverage closest to the nominal level among the compared methods.

Table 2: Results of comparison methods.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model, $n = 400$									
HulC	-	-	-	-	-	-	-	94.8%	0.71(0.21)
BNN	0.06(0.61)	0.50(0.35)	0.02(0.14)	0.12(0.08)	0.27	-	-	97.8%	0.68(0.13)
Ensemble	0.06(0.84)	0.67(0.52)	0.01(0.16)	0.13(0.09)	0.80	0.90	-	95.5%	0.59(0.15)
Naive bp	0.05(0.81)	0.64(0.50)	0.01(0.16)	0.13(0.09)	0.77	0.85	-	95.4%	0.57(0.15)
ESM	0.07(0.50)	0.40(0.31)	0.01(0.10)	0.08(0.06)	0.48	0.59	0.50	94.6%	0.39(0.09)
Logistic Model, $n = 700$									
HulC	-	-	-	-	-	-	-	94.9%	0.67(0.21)
BNN	0.01(0.64)	0.55(0.34)	0.00(0.15)	0.13(0.07)	0.11	-	-	65.7%	0.31(0.10)
Ensemble	0.04(0.48)	0.38(0.30)	0.01(0.10)	0.08(0.06)	0.47	0.53	-	96.8%	0.41(0.10)
Naive bp	0.03(0.48)	0.37(0.29)	0.01(0.10)	0.08(0.06)	0.46	0.52	-	96.7%	0.40(0.09)
ESM	0.06(0.33)	0.26(0.21)	0.01(0.07)	0.06(0.05)	0.31	0.40	0.31	92.9%	0.26(0.06)
Poisson Model, $n = 400$									
HulC	-	-	-	-	-	-	-	90.0%	1.40(0.89)
BNN	0.15(0.49)	0.41(0.30)	0.03(0.36)	0.30(0.21)	0.20	-	-	86.3%	1.55(1.18)
Ensemble	-0.22(0.44)	0.37(0.31)	-0.09(0.30)	0.24(0.21)	0.42	0.50	-	96.0%	1.33(0.71)
Naive bp	-0.21(0.44)	0.37(0.31)	-0.08(0.30)	0.24(0.21)	0.42	0.48	-	95.3%	1.28(0.69)
ESM	-0.14(0.34)	0.29(0.24)	-0.06(0.25)	0.19(0.18)	0.33	0.42	0.35	94.3%	1.02(0.59)
Poisson Model, $n = 700$									
HulC	-	-	-	-	-	-	-	92.8%	1.37(0.84)
BNN	0.26(0.43)	0.40(0.31)	0.12(0.31)	0.28(0.18)	0.09	-	-	73.2%	0.92(0.34)
Ensemble	-0.10(0.30)	0.25(0.20)	-0.04(0.23)	0.17(0.16)	0.29	0.37	-	97.9%	1.06(0.56)
Naive bp	-0.10(0.31)	0.25(0.20)	-0.04(0.24)	0.18(0.16)	0.30	0.35	-	97.1%	1.03(0.54)
ESM	-0.04(0.24)	0.19(0.15)	-0.02(0.19)	0.14(0.13)	0.23	0.32	0.24	94.5%	0.72(0.38)

So, in some cases, using estimators without bias correction or using bootstrap procedures may also have merit, as the presence of bias-correction terms in finite samples can lead to slight undercoverage.

We also investigate the performance of narrower yet deeper architectures under our method. We design a deeper network with $L = 5$ and $\mathbf{p} = (p, 10, 15, 20, 30, 1)$. As deep networks may cause gradient explosion or vanishing during backpropagation (Glorot & Bengio 2010), we incorporate batch normalization layers after the second and second-to-last layers, aiming to mitigate gradient instability while maintaining the network’s ability to learn complex representations. Table 3 shows that key metrics, including empSD, SE, and SE_c, remain largely consistent with those observed for shallower network architectures.

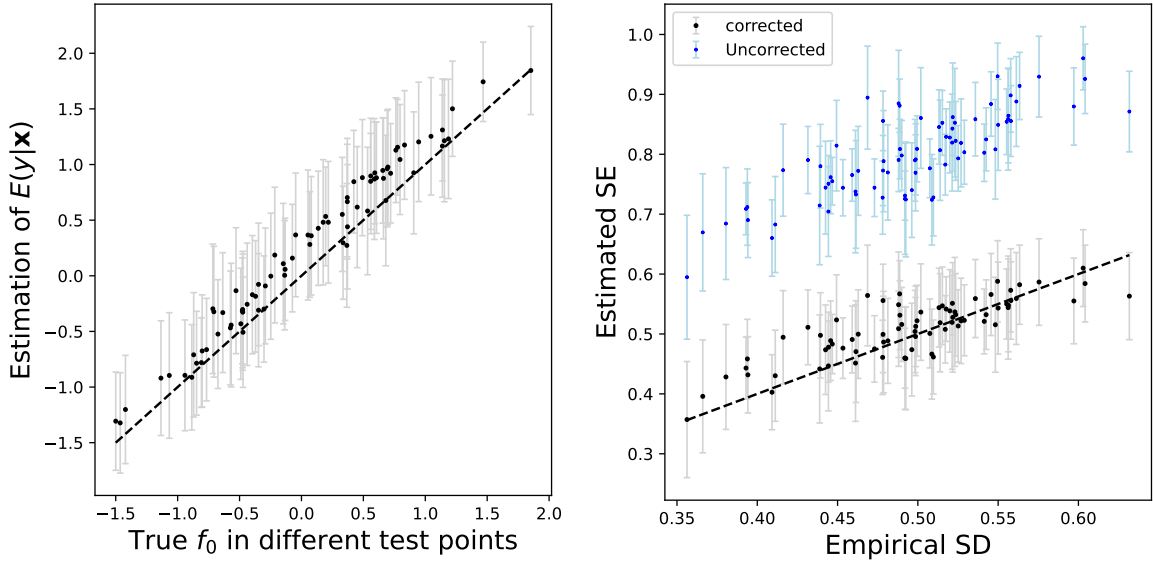
Table 3: Simulation results for different models under varying sample sizes n and subsampling ratios r with different DNNs.

	Bias _{f}	MAE _{f}	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _{c}	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.20(0.51)	0.44(0.33)	0.04(0.11)	0.10(0.07)	0.49	0.80	0.51	90.2%	0.40(0.11)
$r = n^{0.85}$	0.20(0.52)	0.44(0.34)	0.04(0.11)	0.09(0.07)	0.50	0.86	0.52	90.2%	0.41(0.11)
$r = n^{0.9}$	0.19(0.53)	0.45(0.34)	0.04(0.11)	0.10(0.07)	0.51	0.95	0.53	91.0%	0.41(0.11)
$n = 700, r = n^{0.8}$	0.16(0.39)	0.34(0.26)	0.04(0.09)	0.08(0.06)	0.37	0.84	0.38	90.5%	0.31(0.08)
$r = n^{0.85}$	0.20(0.39)	0.35(0.27)	0.04(0.09)	0.08(0.06)	0.37	0.78	0.38	91.1%	0.31(0.08)
$r = n^{0.9}$	0.16(0.39)	0.34(0.26)	0.04(0.09)	0.07(0.06)	0.37	0.84	0.38	91.4%	0.31(0.08)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.35(0.38)	0.38(0.28)	-0.19(0.21)	0.23(0.18)	0.30	0.50	0.31	86.1%	0.70(0.35)
$r = n^{0.85}$	-0.26(0.33)	0.33(0.26)	-0.14(0.21)	0.20(0.16)	0.31	0.55	0.32	90.0%	0.78(0.39)
$r = n^{0.9}$	-0.17(0.32)	0.28(0.23)	-0.08(0.23)	0.18(0.16)	0.31	0.56	0.32	91.8%	0.86(0.45)
$n = 700, r = n^{0.8}$	-0.22(0.24)	0.26(0.20)	-0.13(0.17)	0.16(0.14)	0.23	0.47	0.23	88.0%	0.57(0.27)
$r = n^{0.85}$	-0.13(0.25)	0.21(0.18)	-0.07(0.18)	0.14(0.13)	0.23	0.51	0.23	92.3%	0.63(0.29)
$r = n^{0.9}$	-0.05(0.25)	0.20(0.16)	-0.02(0.19)	0.14(0.13)	0.23	0.56	0.24	93.4%	0.68(0.34)

Nevertheless, coverage probabilities drop to around 91%. Figure 3 presents representative results for $n = 400$ and $r = n^{0.8}$. As shown in Figure 3a, the fitted values exhibit a slight systematic deviation from the diagonal, indicating bias relative to the true function f_0 . Nevertheless, Figure 3b confirms that the corrected variance estimates remain accurate. This bias shifts the confidence intervals, leading to reduced coverage despite precise variance estimation. Similar bias-related undercoverage has been reported for random forests in Wager & Athey (2018). In our case, greater network depth may induce complex interactions that, without sufficient data, amplify estimation bias.

5 Analysis of eICU data

Preventing ICU readmissions is vital as readmitted patients face higher mortality risks and increased expenses. Identifying high-risk individuals enables more personalized care and optimized resource use. The eICU Collaborative Research Database (eICU-CRD)



(a) Point Estimates with Intervals (b) Estimated Standard Errors and Variations

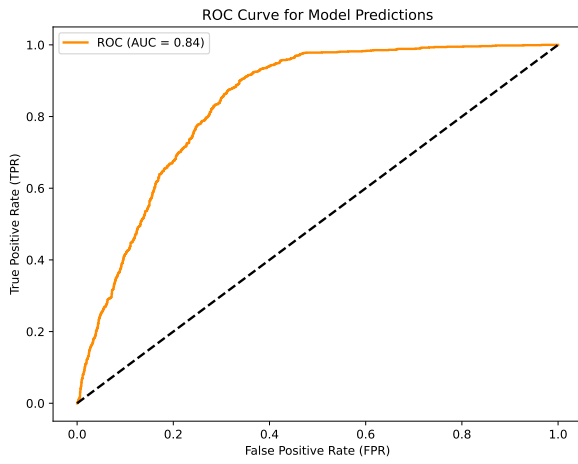
Figure 3: Estimation and inference in simulation samples: Logistic Model with $n = 400$ and $r = n^{0.8}$ under a deeper network. Figure 3a displays the average estimated $\mathbb{E}(y|\mathbf{x})$ and its deviation (gray band) across 300 runs at various test points. Figure 3b compares corrected and uncorrected SEs and their deviations (gray and blue bands) across 300 runs with empirical SDs across test samples.

(Pollard et al. 2018) is a large-scale, multi-center collection of anonymized ICU patient data, created through collaboration between the MIT Laboratory for Computational Physiology and healthcare organizations. It includes detailed information on patients’ physiological measurements, clinical data, lab results, medication usage, and outcomes during their ICU stays. Our analysis focuses on Hospital #188, a major hospital with 2,632 patients admitted to ICU during 2014–2015. We include ten covariates in our analysis, selected for their clinical relevance and frequent use in prior predictive modeling studies. The features include six lab tests (Hct, chloride, WBC, Hgb, RBC, glucose), patient age, Emergency Department admission status, and two APACHE-based severity indicators: anticipated hospital stay

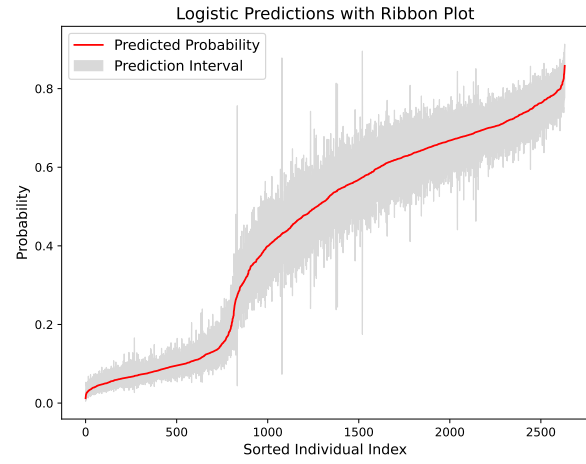
and mortality likelihood. Among the 2,632 patients admitted to the ICU, 1,476 were not readmitted, 1,046 were readmitted once, 77 were readmitted twice, 26 were readmitted three times, and 7 were readmitted five or more times. We apply nonparametric logistic and Poisson models within our ESM framework to estimate, respectively, the probability of ICU readmission and total number of ICU readmissions using the laboratory and demographic information. We validate our results using 5-fold cross-validation. The dataset is split into five equal-sized folds; the model is trained on four folds ($n = 2105$), and the ESM method is applied to estimate the subject-specific means and confidence intervals for each subject in the held-out fold, using only covariates while withholding true outcomes. Repeating this process across all folds yields point estimates and inference intervals for all observations.

To estimate the probability of ICU readmission, we define a binary outcome variable: $y = 1$ if any readmission occurs during follow-up, i.e., 2014–2015, and $y = 0$ otherwise. When implementing our method, we set the subsampling size to $r = n^{0.9} = 979$ and the number of subsampling iterations to $B = 3000$. We use a (10, 128, 64, 1) architecture with ReLU activation, selected for its balance of accuracy and efficiency via simulations. Training is performed over 300 epochs with a tuned learning rate of $\eta = 0.1$.

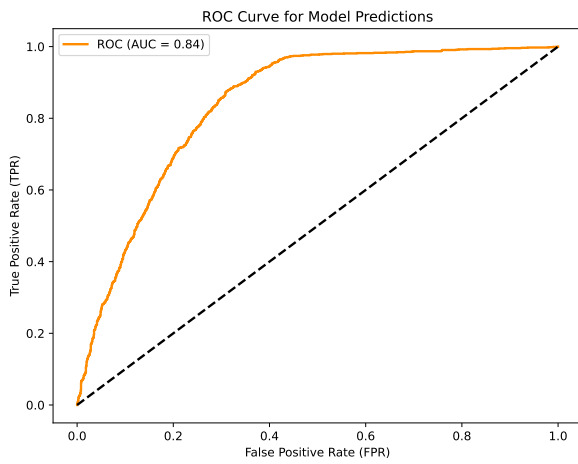
As shown in Figure 4, both random forests and neural networks achieve similar predictive performance (AUC = 0.84), but their uncertainty estimates differ much. Neural networks yield shorter average confidence intervals (0.123 vs. 0.225 for RFs). Figures 4b and 4d highlight these differences by displaying predicted readmission probabilities with 95% intervals across patients sorted by risk. The variability in interval widths suggests heteroskedasticity, likely driven by differences in data quality or patient characteristics. Notably, Figure 4b reveals a subset of patients (indices between 800–1600) with unusually high uncertainty under DNNs; further details are provided in Section B.1.




(a) ROC curve in DNNs



(b) Individual estimates and CIs in DNNs




(c) ROC curve in RFs

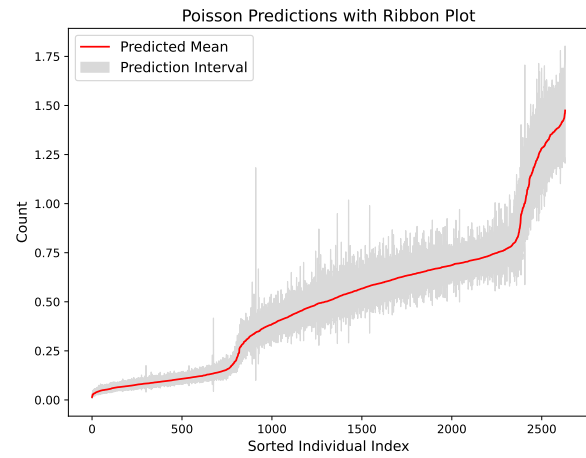


(d) Individual estimates and CIs in RFs

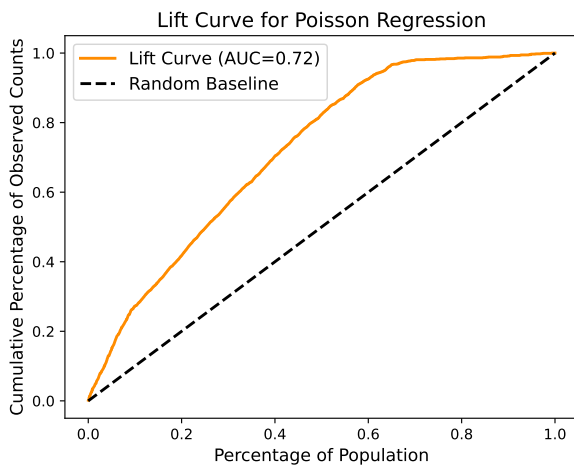
Figure 4: Evaluation of the nonparametric logistic model estimator. Figure 4a and 4c show the ROC curve with both AUC of 0.84. Figure 4b and 4d display estimated subject-level probabilities of ICU readmission and confidence intervals, illustrating heteroskedasticity across individuals.



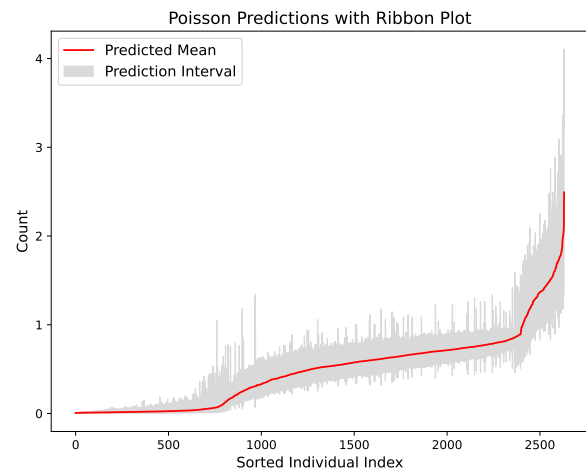
(a) Lift curve in DNNs



(b) Individual estimates and CIs in DNNs



(c) Lift curve in RFs



(d) Individual estimates and CIs in RFs

Figure 5: Evaluation of the nonparametric logistic model estimator. Figure 5a and 5c show the ROC curve with both AUC of 0.72. Figure 5b and 5d display estimated subject-level probabilities of ICU readmission and confidence intervals, illustrating heteroskedasticity across individuals.

To estimate subject-level ICU readmissions during 2014–2015, we fit a nonparametric Poisson regression using a DNN to model $\mathbb{E}(y|\mathbf{x}_*)$, offering a complementary view of patient well-being. Figure 5 shows a lift curve around 0.72, indicating comparable predictive accuracy for DNNs and RFs. However, uncertainty estimates differ: DNNs yield average interval lengths of 0.124, compared to 0.274 for RFs. Figures 5b and 5d illustrate individual predictions with 95% intervals, again showing heteroskedasticity. Additional analyses (Section B.2) reveal that in some settings, RFs produce narrower intervals than DNNs, highlighting model sensitivity to data structure and signal-to-noise characteristics. We also assess the transferability of ESM by applying models trained on hospital #188 directly to data from hospital #458; see Section B.3 in the Supplementary Material. The results indicate that the model remains transferable in this setting.

6 Discussion

This paper presents a general framework for inference on subject-level means estimated by deep neural networks for categorical and exponential family outcomes. We construct a DNN estimator by minimizing the loss function induced from the generalized nonparametric regression model (GNRM). To address a key gap in leveraging this general loss function, we analyze the convergence rates of DNNs under the GNRM framework and establish connections with U-statistics and Hoeffding decomposition theory. Building on these results, we introduce an ensemble subsampling method (ESM) to enable valid inference. Numerical studies and an application to the eICU dataset demonstrate ESM’s utility in quantifying predictive uncertainty and supporting personalized care.

Our methodology may have broad applicability as it enables valid inference for any continuously-differentiable functional of $f_0(x)$, such as $\text{Var}(y|\mathbf{x}) = \psi''(f_0(x)) =: g(f_0(x))$,

where $g(u) = \psi''(u)$, the second derivative of $\psi(u)$. We can estimate $\text{Var}(y|\mathbf{x})$ via the plug-in estimator $g(\widehat{f}^B(\mathbf{x}))$ and apply the delta method to give $\sqrt{\frac{n}{r^2\xi_{1,r}(\mathbf{x})}} \cdot (g(\widehat{f}^B(\mathbf{x})) - g(f_0(\mathbf{x}))) = \sqrt{\frac{n}{r^2\xi_{1,r}(\mathbf{x})}} \cdot g'(f_0(\mathbf{x}))(\widehat{f}^B(\mathbf{x}) - f_0(\mathbf{x})) + o_p(1)$, which will lead to estimated confidence intervals. Our method may also naturally extend to semiparametric causal inference settings where the target parameter can be expressed as a functional of the conditional means. Specifically, to estimate the Conditional Average Treatment Effect (CATE), $\tau(\mathbf{x}) = \mathbb{E}[y(1)|\mathbf{x}] - \mathbb{E}[y(0)|\mathbf{x}]$, we apply ESM separately to treated and control groups under standard unconfoundedness. Multiple subsamples are drawn for each group, predictive models are trained to estimate $\mathbb{E}[y(g)|\mathbf{x}]$, and ensemble means $\widehat{\mu}_1(\mathbf{x})$ and $\widehat{\mu}_0(\mathbf{x})$ are obtained. The CATE estimator is $\widehat{\tau}(\mathbf{x}) = \widehat{\mu}_1(\mathbf{x}) - \widehat{\mu}_0(\mathbf{x})$. ESM’s variance framework provides inference by obtaining $\widehat{\sigma}_g^2(\mathbf{x}) = \widehat{\text{Var}}\{\widehat{\mu}_g(\mathbf{x})\}$ for $g = 0, 1$, yielding pointwise confidence intervals $\widehat{\tau}(\mathbf{x}) \pm z_{1-\alpha/2}\sqrt{\widehat{\sigma}_1^2(\mathbf{x}) + \widehat{\sigma}_0^2(\mathbf{x})}$. Detailed justifications warrant investigation.

Our theoretical analysis focuses on ReLU activations, where monomial-based approximations are tractable. We envision that ESM can be extended to networks with smooth activation functions such as tanh or GELU, although this requires careful consideration. Unlike ReLU, smooth activations are less flexible in local adaptations, which may result in slower approximation rates and influence convergence behavior. To preserve inference validity in this setting, it may be necessary to employ alternative analytical tools or apply stronger regularization to control model complexity and prevent saturation. Moreover, Our theoretical analysis relies on several assumptions that are broadly consistent with recent work on implicit bias and convergence in overparameterized models ([Chizat & Bach 2020](#), [Chatterji et al. 2021](#), [Cao et al. 2022](#)). We acknowledge, however, that some of these assumptions may appear strong. In particular, the conditions imposed on $\xi_{1,r}(\mathbf{x})$, which ensure the validity of the first-order Hoeffding projection, can be further improved. An im-

portant direction for future work is to relax these conditions by linking them more directly to of the neural network class, such as its sparsity or other structural characteristics.

Data Availability Statement. The data used in this study are publicly available from the eICU Collaborative Research Database (<https://eicu-crd.mit.edu/>) upon completion of a data use agreement and approval through the PhysioNet credentialed access process.

Acknowledgements. We are deeply grateful to the Editor, Associate Editor, two referees and one reproducibility reviewer for insightful reviews that have greatly improved the manuscript.

References

- Alaa, A. & Van Der Schaar, M. (2020), Discriminative jackknife: Quantifying uncertainty in deep learning via higher-order influence functions, *in* ‘International Conference on Machine Learning’, pp. 165–174.
- Angelopoulos, A. N. & Bates, S. (2021), ‘A gentle introduction to conformal prediction and distribution-free uncertainty quantification’, *arXiv preprint arXiv:2107.07511* .
- Athey, S., Imbens, G. W. & Wager, S. (2018), ‘Approximate residual balancing: debiased inference of average treatment effects in high dimensions’, *Journal of the Royal Statistical Society Series B: Statistical Methodology* **80**(4), 597–623.
- Barber, R. F., Candes, E. J., Ramdas, A. & Tibshirani, R. J. (2021), ‘Predictive inference with the jackknife+’, *The Annals of Statistics* **49**(1), 486–507.
- Bartlett, P. L., Long, P. M., Lugosi, G. & Tsigler, A. (2020), ‘Benign overfitting in linear regression’, *Proceedings of the National Academy of Sciences* **117**(48), 30063–30070.
- Bhattacharya, S., Fan, J. & Mukherjee, D. (2024), ‘Deep neural networks for nonparametric interaction models with diverging dimension’, *The Annals of Statistics* **52**(6), 2738–2766.
- Boroskikh, Y. V. (2020), *U-statistics in Banach Spaces*, Walter de Gruyter GmbH & Co KG.
- Breiman, L. (2001), ‘Random forests’, *Machine Learning* **45**(1), 5–32.
- Brown, L. D. (1986), ‘Fundamentals of statistical exponential families: with applications in statistical decision theory’, *Lecture Notes-Monograph Series* **9**, i–279.

- Cao, Y., Chen, Z., Belkin, M. & Gu, Q. (2022), ‘Benign overfitting in two-layer convolutional neural networks’, *Advances in Neural Information Processing Systems* **35**, 25237–25250.
- Chatterji, N. S., Long, P. M. & Bartlett, P. L. (2021), ‘When does gradient descent with logistic loss find interpolating two-layer networks?’, *Journal of Machine Learning Research* **22**(159), 1–48.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. & Robins, J. (2018), ‘Double/debiased machine learning for treatment and structural parameters’, *The Econometrics Journal* **21**(1), C1–C68.
- Chizat, L. & Bach, F. (2020), Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss, in ‘Conference on learning theory’, pp. 1305–1338.
- Drucker, H., Burges, C. J., Kaufman, L., Smola, A. & Vapnik, V. (1996), ‘Support vector regression machines’, *Advances in Neural Information Processing Systems* **9**, 155–161.
- Fan, J. (2018), *Local Polynomial Modelling and Its Applications: Monographs on Statistics and Applied Probability 66*, Routledge.
- Fan, J. & Gu, Y. (2024), ‘Factor augmented sparse throughput deep ReLU neural networks for high dimensional regression’, *Journal of the American Statistical Association* **119**(548), 2680–2694.
- Fan, J., Gu, Y. & Zhou, W.-X. (2024), ‘How do noise tails impact on deep ReLU networks?’, *The Annals of Statistics* **52**(4), 1845–1871.
- Fei, Z. & Li, Y. (2021), ‘Estimation and inference for high dimensional generalized linear models: A splitting and smoothing approach’, *Journal of Machine Learning Research* **22**(58), 1–32.
- Fei, Z. & Li, Y. (2024), ‘U-learning for prediction inference via combinatorial multi-subsampling: With applications to lasso and neural networks’, *arXiv preprint arXiv:2407.15301*.
- Frees, E. W. (1989), ‘Infinite order U-statistics’, *Scandinavian Journal of Statistics* **16**(1), 29–45.
- Glorot, X. & Bengio, Y. (2010), Understanding the difficulty of training deep feedforward neural networks, in ‘Proceedings of the thirteenth international conference on artificial intelligence and statistics’, JMLR Workshop and Conference Proceedings, pp. 249–256.

- Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. (2016), *Deep Learning*, Vol. 1, MIT press Cambridge.
- Guo, Z., Rakshit, P., Herman, D. S. & Chen, J. (2021), ‘Inference for the case probability in high-dimensional logistic regression’, *Journal of Machine Learning Research* **22**(254), 1–54.
- Györfi, L., Kohler, M., Krzyżak, A. & Walk, H. (2006), *A Distribution-Free Theory of Nonparametric Regression*, Springer Science & Business Media.
- Hastie, T., Tibshirani, R., Friedman, J. H. & Friedman, J. H. (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer.
- Hoeffding, W. (1992), ‘A class of statistics with asymptotically normal distribution’, *Breakthroughs in Statistics: Foundations and Basic Theory* pp. 308–334.
- Hornik, K., Stinchcombe, M. & White, H. (1989), ‘Multilayer feedforward networks are universal approximators’, *Neural Networks* **2**(5), 359–366.
- Huang, J., Xi, H., Zhang, L., Yao, H., Qiu, Y. & Wei, H. (2024), Conformal prediction for deep classifier via label ranking, in ‘International Conference on Machine Learning’, pp. 20331–20347.
- Jiao, Y., Shen, G., Lin, Y. & Huang, J. (2023), ‘Deep nonparametric regression on approximate manifolds: Nonasymptotic error bounds with polynomial prefactors’, *The Annals of Statistics* **51**(2), 691–716.
- Kim, B., Xu, C. & Barber, R. (2020), ‘Predictive inference is free with the jackknife+–after-bootstrap’, *Advances in Neural Information Processing Systems* **33**, 4138–4149.
- Kohler, M. & Krzyżak, A. (2005), ‘Adaptive regression estimation with multilayer feedforward neural networks’, *Nonparametric Statistics* **17**(8), 891–913.
- Kohler, M. & Langer, S. (2021), ‘On the rate of convergence of fully connected deep neural network regression estimates’, *The Annals of Statistics* **49**(4), 2231–2249.
- Kou, Y., Chen, Z., Chen, Y. & Gu, Q. (2023), Benign overfitting in two-layer ReLU convolutional neural networks, in ‘International Conference on Machine Learning’, pp. 17615–17659.
- Kuchibhotla, A. K., Balakrishnan, S. & Wasserman, L. (2024), ‘The HulC: Confidence regions from convex hulls’, *Journal of the Royal Statistical Society Series B: Statistical Methodology* **86**(3), 586–622.

- Lakshminarayanan, B., Pritzel, A. & Blundell, C. (2017), ‘Simple and scalable predictive uncertainty estimation using deep ensembles’, *Advances in Neural Information Processing Systems* **30**.
- LeCun, Y., Bengio, Y. & Hinton, G. (2015), ‘Deep learning’, *Nature* **521**(7553), 436–444.
- Lee, A. J. (2019), *U-statistics: Theory and Practice*, Routledge.
- Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R. J. & Wasserman, L. (2018), ‘Distribution-free predictive inference for regression’, *Journal of the American Statistical Association* **113**(523), 1094–1111.
- Lu, J., Shen, Z., Yang, H. & Zhang, S. (2021), ‘Deep network approximation for smooth functions’, *SIAM Journal on Mathematical Analysis* **53**(5), 5465–5506.
- Lundberg, S. M., Nair, B., Vavilala, M. S., Horibe, M., Eisses, M. J., Adams, T., Liston, D. E., Low, D. K.-W., Newman, S.-F., Kim, J. et al. (2018), ‘Explainable machine-learning predictions for the prevention of hypoxaemia during surgery’, *Nature Biomedical Engineering* **2**(10), 749–760.
- McCaffrey, D. F. & Gallant, A. R. (1994), ‘Convergence rates for single hidden layer feed-forward networks’, *Neural Networks* **7**(1), 147–158.
- Mentch, L. & Hooker, G. (2016), ‘Quantifying uncertainty in random forests via confidence intervals and hypothesis tests’, *Journal of Machine Learning Research* **17**(26), 1–41.
- Nadaraya, E. A. (1964), ‘On estimating regression’, *Theory of Probability & Its Applications* **9**(1), 141–142.
- Nair, V. & Hinton, G. E. (2010), Rectified linear units improve restricted boltzmann machines, in ‘International Conference on Machine Learning’, pp. 807–814.
- Peng, W., Mentch, L. & Stefanski, L. (2024), ‘Bias, consistency, and alternative perspectives of the infinitesimal jackknife’, *Statistica Sinica* .
- Pollard, T. J., Johnson, A. E. W., Raffa, J. D., Celi, L. A., Mark, R. G. & Badawi, O. (2018), ‘The eICU collaborative research database, a freely available multi-center database for critical care research’, *Scientific Data* **5**(1), 1–13.
- Schmidt-Hieber, A. J. (2020), ‘Nonparametric regression using deep neural networks with ReLU activation function’, *The Annals of Statistics* **48**(4), 1875–1897.
- Schupbach, J., Sheppard, J. W. & Forrester, T. (2020), Quantifying uncertainty in neural network ensembles using U-statistics, in ‘2020 International Joint Conference on Neural Networks (IJCNN)’, IEEE, pp. 1–8.

- Shen, G., Jiao, Y., Lin, Y. & Huang, J. (2021), ‘Robust nonparametric regression with deep neural networks’, *arXiv preprint arXiv:2107.10343* .
- Shen, Z., Yang, H. & Zhang, S. (2019), ‘Nonlinear approximation via compositions’, *Neural Networks* **119**, 74–84.
- Sun, Y., Xiong, W. & Liang, F. (2021), ‘Sparse deep learning: A new framework immune to local traps and miscalibration’, *Advances in Neural Information Processing Systems* **34**, 22301–22312.
- Takezawa, K. (2005), *Introduction to Nonparametric Regression*, John Wiley & Sons.
- Van Calster, B., McLernon, D. J., Van Smeden, M., Wynants, L., Steyerberg, E. W., diagnostic tests, T. G. E. & prediction models’ of the STRATOS initiative Bossuyt Patrick Collins Gary S. Macaskill Petra McLernon David J. Moons Karel GM Steyerberg Ewout W. Van Calster Ben van Smeden Maarten Vickers Andrew J. (2019), ‘Calibration: the achilles heel of predictive analytics’, *BMC Medicine* **17**(1), 230.
- Wager, S. & Athey, S. (2018), ‘Estimation and inference of heterogeneous treatment effects using random forests’, *Journal of the American Statistical Association* **113**(523), 1228–1242.
- Wager, S., Hastie, T. & Efron, B. (2014), ‘Confidence intervals for random forests: The jackknife and the infinitesimal jackknife’, *Journal of Machine Learning Research* **15**(1), 1625–1651.
- Wang, Q. & Wei, Y. (2022), ‘Quantifying uncertainty of subsampling-based ensemble methods under a U-statistic framework’, *Journal of Statistical Computation and Simulation* **92**(17), 3706–3726.
- Wang, X., Zhou, L. & Lin, H. (2024), ‘Deep regression learning with optimal loss function’, *Journal of the American Statistical Association* pp. 1–13.
- Yan, S., Yao, F. & Zhou, H. (2025), ‘Deep regression for repeated measurements’, *Journal of the American Statistical Association* pp. 1–12.
- Yarotsky, D. (2017), ‘Error bounds for approximations with deep ReLU networks’, *Neural Networks* **94**, 103–114.
- Zhang, Y. & Politis, D. N. (2023), ‘Bootstrap prediction intervals with asymptotic conditional validity and unconditional guarantees’, *Information and Inference: A Journal of the IMA* **12**(1), 157–209.

Supplementary Material for “Inference for Deep Neural Network Estimators in Generalized Nonparametric Models”

A Additional Numerical Results

We conduct additional numerical experiments to further evaluate our proposed method. In Section A.1, we extend the approach to the binomial regression setting. In Section A.2, we compare neural networks and random forests, showing that their relative performance varies across different scenarios. In Section A.3, we further examine simpler kernel regression models to assess ESM under smoother nonlinear structures. In Section A.4, we present experiments illustrating how our method performs as data dimensionality and nonlinearity increase. We also investigate the effects of key optimization and design parameters, including the subsample size r , the number of subsamples B , the network depth L , and the dropout rate, in Sections A.5, A.6, A.7, A.8, and A.9, respectively.

A.1 Binomial Regression

We show the results under binomial regression, that is, with $k \in \{0, 1, \dots, n_{\text{trial}}\}$,

$$\mathbb{P}(y_i = k | \mathbf{x}_i) = \binom{n_{\text{trial}}}{k} p^k(\mathbf{x}^{(i)}) (1 - p(\mathbf{x}^{(i)}))^k,$$

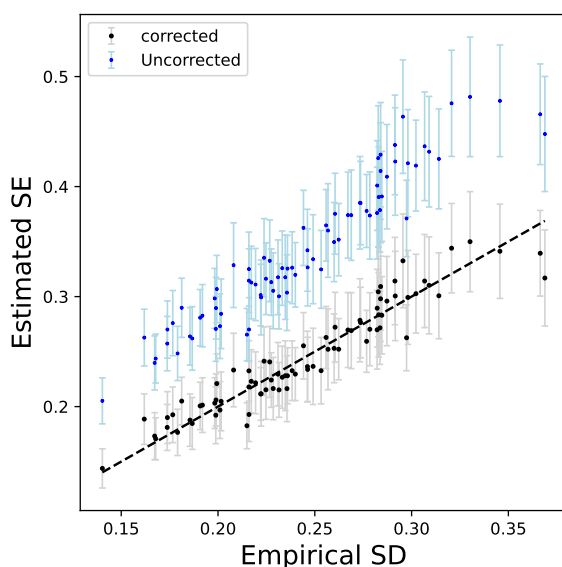
with $p(\mathbf{x}^{(i)}) = \frac{1}{1 + \exp\{-g(\mathbf{x}_i)\}}$. Under this model, $f_0(\mathbf{x}) = g(\mathbf{x})$ and $\psi'(x) = n_{\text{trial}}(1 + \exp(-x))^{-1}$.

We use the same simulation framework as described in Section 4. We set the number of trials in the binomial distribution to $n_{\text{trial}} = 5$, while keeping all other simulation parameters and settings consistent with the baseline experiments. At the same time, we also vary the resampling number B from 1400 to 3000, and present the corresponding results to examine


the effect of B on the variance estimation and interval construction.

Table S.1: Simulation results for Binomial regression under varying sample sizes n and subsampling ratios r .

	Bias _{f}	MAE _{f}	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _{c}	CP	AIL
$B = 1400$									
$n = 400, r = n^{0.8}$	0.03(0.34)	0.27(0.21)	0.01(0.07)	0.06(0.04)	0.31	0.39	0.32	94.0%	0.26(0.07)
$r = n^{0.85}$	0.04(0.36)	0.28(0.23)	0.01(0.07)	0.06(0.05)	0.34	0.41	0.34	93.5%	0.27(0.08)
$r = n^{0.9}$	0.04(0.34)	0.27(0.21)	0.01(0.07)	0.06(0.04)	0.32	0.43	0.35	95.0%	0.28(0.08)
$n = 700, r = n^{0.8}$	0.02(0.27)	0.21(0.17)	0.00(0.06)	0.05(0.03)	0.25	0.34	0.25	93.2%	0.21(0.05)
$r = n^{0.85}$	0.03(0.25)	0.20(0.16)	0.01(0.05)	0.04(0.03)	0.24	0.34	0.25	94.3%	0.21(0.05)
$r = n^{0.9}$	0.03(0.26)	0.20(0.16)	0.01(0.06)	0.04(0.03)	0.25	0.34	0.25	93.5%	0.21(0.05)
$B = 3000$									
$n = 400, r = n^{0.8}$	0.05(0.34)	0.27(0.21)	0.01(0.07)	0.06(0.04)	0.31	0.36	0.32	93.9%	0.26(0.07)
$r = n^{0.85}$	0.03(0.35)	0.27(0.22)	0.01(0.07)	0.06(0.05)	0.33	0.38	0.34	94.3%	0.27(0.07)
$r = n^{0.9}$	0.03(0.34)	0.27(0.21)	0.00(0.07)	0.06(0.04)	0.32	0.38	0.34	95.0%	0.28(0.07)
$n = 700, r = n^{0.8}$	0.04(0.26)	0.20(0.16)	0.01(0.06)	0.04(0.03)	0.24	0.29	0.25	93.3%	0.20(0.05)
$r = n^{0.85}$	0.01(0.25)	0.19(0.15)	0.00(0.05)	0.04(0.03)	0.23	0.29	0.24	94.4%	0.20(0.05)
$r = n^{0.9}$	0.02(0.24)	0.19(0.15)	0.00(0.05)	0.04(0.03)	0.23	0.29	0.24	94.1%	0.20(0.05)



(a) Estimated Standard Errors



(b) Estimated Standard Errors

Figure S.1: Estimation and inference in simulation samples: Binomial Model with $n = 700$ and $r = n^{0.9}$. Figure S.1a and S.1b show corrected and uncorrected estimation SE versus empirical SD of all test samples with $B = 1400$ and $B = 3000$, respectively.

As shown in Table S.1, our method achieves reasonable coverage probabilities and stable interval lengths under the binomial regression setting, across various values of B and r . In binomial regression, the corrected standard errors SE_c closely track the empirical variability, supporting the validity of our variance estimation procedure. We also plot the corrected and uncorrected estimation SE versus empirical SD with different B in Figure S.1. Figure S.1 also shows that with a larger B , the uncorrected estimator becomes more accurate and gradually approaches the empirical SD. In contrast, the corrected SE consistently aligns well with the empirical SD across different values of B , demonstrating the effectiveness of our correction procedure, especially in finite-sample settings.

A.2 Comparison to Random Forests

Although structurally different from deep neural networks (DNNs), random forests (RFs) are another widely used tool for nonparametric estimation. As a benchmark, we compare the performance of neural networks and random forests within our simulation framework to examine their respective behaviors across varying settings, rather than to identify a universally superior method. In addition, we investigate the use of our corrected variance estimator within the RF framework to enhance the accuracy of empirical standard deviation estimation.

Table S.2: Comparison under the basic setting in logistic model.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model with DNNs									
$n = 400, r = n^{0.8}$	0.11(0.72)	0.57(0.44)	0.02(0.14)	0.11(0.08)	0.65	0.78	0.66	91.7%	0.46(0.13)
$r = n^{0.85}$	0.06(0.59)	0.47(0.36)	0.01(0.12)	0.10(0.07)	0.55	0.68	0.58	94.0%	0.43(0.11)
$r = n^{0.9}$	0.07(0.50)	0.40(0.31)	0.01(0.10)	0.08(0.06)	0.48	0.59	0.50	94.6%	0.39(0.09)
$n = 700, r = n^{0.8}$	0.08(0.41)	0.33(0.26)	0.02(0.09)	0.07(0.05)	0.39	0.51	0.39	93.6%	0.31(0.08)
$r = n^{0.85}$	0.05(0.36)	0.28(0.22)	0.01(0.08)	0.06(0.05)	0.34	0.44	0.34	93.9%	0.28(0.06)
$r = n^{0.9}$	0.06(0.33)	0.26(0.21)	0.01(0.07)	0.06(0.05)	0.31	0.40	0.31	93.5%	0.26(0.06)
Logistic Model with RFs									
$n = 400, r = n^{0.8}$	0.07(0.39)	0.31(0.24)	0.02(0.09)	0.07(0.05)	0.31	0.38	0.33	89.1%	0.28(0.03)
$r = n^{0.85}$	0.07(0.41)	0.33(0.25)	0.02(0.09)	0.07(0.06)	0.36	0.43	0.36	91.1%	0.31(0.04)
$r = n^{0.9}$	0.08(0.44)	0.35(0.27)	0.02(0.10)	0.08(0.06)	0.40	0.48	0.41	93.0%	0.35(0.05)
$n = 700, r = n^{0.8}$	0.07(0.33)	0.27(0.21)	0.02(0.07)	0.06(0.05)	0.28	0.36	0.28	89.1%	0.25(0.03)
$r = n^{0.85}$	0.06(0.36)	0.29(0.22)	0.01(0.08)	0.06(0.05)	0.32	0.41	0.32	91.8%	0.27(0.03)
$r = n^{0.9}$	0.06(0.37)	0.30(0.23)	0.01(0.08)	0.07(0.05)	0.35	0.47	0.37	94.2%	0.31(0.04)

Table S.3: Comparison under the basic setting in Poisson model.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Poisson Model with DNNs									
$n = 400, r = n^{0.8}$	-0.32(0.38)	0.38(0.31)	-0.16(0.23)	0.22(0.18)	0.35	0.43	0.36	88.6%	0.86(0.45)
$r = n^{0.85}$	-0.23(0.36)	0.33(0.27)	-0.12(0.25)	0.21(0.18)	0.34	0.43	0.36	91.9%	0.95(0.52)
$r = n^{0.9}$	-0.14(0.34)	0.29(0.24)	-0.06(0.25)	0.19(0.18)	0.33	0.42	0.35	94.3%	1.02(0.59)
$n = 700, r = n^{0.8}$	-0.17(0.26)	0.25(0.20)	-0.10(0.19)	0.16(0.14)	0.25	0.34	0.26	90.7%	0.67(0.34)
$r = n^{0.85}$	-0.10(0.26)	0.21(0.17)	-0.06(0.19)	0.15(0.14)	0.24	0.33	0.25	93.2%	0.71(0.38)
$r = n^{0.9}$	-0.04(0.24)	0.19(0.15)	-0.02(0.19)	0.14(0.13)	0.23	0.32	0.24	94.5%	0.72(0.38)
Poisson Model with RFs									
$n = 400, r = n^{0.8}$	-0.03(0.24)	0.19(0.16)	-0.01(0.17)	0.13(0.11)	0.21	0.25	0.21	92.4%	0.60(0.21)
$r = n^{0.85}$	-0.03(0.26)	0.20(0.17)	0.00(0.18)	0.14(0.11)	0.24	0.29	0.24	93.7%	0.68(0.25)
$r = n^{0.9}$	-0.02(0.27)	0.21(0.18)	0.01(0.20)	0.15(0.13)	0.26	0.32	0.28	94.6%	0.79(0.30)
$n = 700, r = n^{0.8}$	-0.03(0.21)	0.16(0.14)	-0.01(0.15)	0.11(0.09)	0.18	0.24	0.19	92.8%	0.52(0.17)
$r = n^{0.85}$	-0.03(0.22)	0.17(0.14)	0.00(0.16)	0.12(0.10)	0.20	0.28	0.21	94.2%	0.59(0.20)
$r = n^{0.9}$	-0.03(0.26)	0.20(0.17)	0.00(0.18)	0.14(0.12)	0.24	0.32	0.25	94.3%	0.69(0.26)

We first present random forest performance under the basic settings as in the simulation section of the main text, covering logistic and Poisson scenarios and using the same evaluation metrics as for neural networks.

Table S.4: Comparison under the changing setting in logistic model.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model with DNNs									
$n = 400, r = n^{0.8}$	0.07(0.45)	0.36(0.28)	0.01(0.09)	0.07(0.06)	0.42	0.50	0.43	93.8%	0.32(0.09)
$r = n^{0.85}$	0.04(0.40)	0.32(0.25)	0.01(0.08)	0.07(0.05)	0.38	0.45	0.38	93.9%	0.30(0.08)
$r = n^{0.9}$	0.04(0.37)	0.29(0.23)	0.01(0.08)	0.06(0.05)	0.34	0.41	0.35	93.8%	0.28(0.07)
$n = 700, r = n^{0.8}$	0.05(0.29)	0.23(0.19)	0.01(0.06)	0.05(0.04)	0.26	0.34	0.27	93.3%	0.22(0.06)
$r = n^{0.85}$	0.04(0.29)	0.23(0.19)	0.01(0.06)	0.05(0.04)	0.25	0.32	0.25	90.4%	0.21(0.05)
$r = n^{0.9}$	0.04(0.29)	0.22(0.19)	0.01(0.06)	0.05(0.04)	0.22	0.31	0.23	90.0%	0.19(0.05)
Logistic Model with RFs									
$n = 400, r = n^{0.8}$	0.02(0.46)	0.36(0.30)	0.01(0.09)	0.07(0.06)	0.32	0.38	0.32	84.9%	0.28(0.04)
$r = n^{0.85}$	0.04(0.46)	0.36(0.30)	0.01(0.10)	0.08(0.06)	0.35	0.42	0.36	87.6%	0.31(0.05)
$r = n^{0.9}$	0.03(0.48)	0.37(0.30)	0.01(0.10)	0.08(0.06)	0.38	0.48	0.41	90.8%	0.34(0.05)
$n = 700, r = n^{0.8}$	0.03(0.40)	0.31(0.26)	0.01(0.08)	0.06(0.05)	0.28	0.36	0.28	84.8%	0.24(0.03)
$r = n^{0.85}$	0.02(0.41)	0.32(0.26)	0.01(0.08)	0.07(0.05)	0.32	0.41	0.32	88.5%	0.27(0.04)
$r = n^{0.9}$	0.01(0.42)	0.33(0.27)	0.01(0.09)	0.07(0.05)	0.34	0.47	0.37	91.3%	0.31(0.05)

Table S.5: Comparison under the changing setting in Poisson model.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Poisson Model with DNNs									
$n = 400, r = n^{0.8}$	-0.20(0.31)	0.28(0.24)	-0.12(0.24)	0.19(0.19)	0.29	0.35	0.30	90.8%	0.80(0.52)
$r = n^{0.85}$	-0.12(0.29)	0.24(0.20)	-0.07(0.24)	0.17(0.18)	0.27	0.34	0.29	93.3%	0.82(0.58)
$r = n^{0.9}$	-0.06(0.27)	0.21(0.18)	-0.03(0.24)	0.16(0.18)	0.25	0.32	0.27	95.0%	0.81(0.59)
$n = 700, r = n^{0.8}$	-0.10(0.22)	0.18(0.15)	-0.07(0.18)	0.13(0.14)	0.20	0.27	0.20	92.3%	0.59(0.37)
$r = n^{0.85}$	-0.04(0.20)	0.16(0.13)	-0.03(0.17)	0.12(0.13)	0.19	0.26	0.19	94.1%	0.59(0.38)
$r = n^{0.9}$	-0.01(0.20)	0.15(0.12)	-0.01(0.17)	0.12(0.12)	0.18	0.24	0.18	93.4%	0.58(0.37)
Poisson Model with RFs									
$n = 400, r = n^{0.8}$	-0.00(0.24)	0.19(0.15)	-0.00(0.20)	0.14(0.14)	0.21	0.25	0.21	91.2%	0.62(0.25)
$r = n^{0.85}$	-0.00(0.26)	0.20(0.16)	0.00(0.20)	0.15(0.14)	0.23	0.28	0.24	93.3%	0.71(0.30)
$r = n^{0.9}$	-0.00(0.28)	0.22(0.18)	0.02(0.23)	0.17(0.16)	0.26	0.33	0.28	93.4%	0.84(0.40)
$n = 700, r = n^{0.8}$	0.00(0.21)	0.16(0.14)	-0.00(0.17)	0.12(0.12)	0.18	0.24	0.18	91.4%	0.54(0.21)
$r = n^{0.85}$	0.01(0.22)	0.17(0.14)	0.01(0.18)	0.13(0.12)	0.20	0.27	0.21	93.2%	0.63(0.26)
$r = n^{0.9}$	-0.00(0.26)	0.20(0.16)	0.01(0.20)	0.15(0.13)	0.24	0.32	0.25	94.1%	0.73(0.33)

As shown in Tables S.2 and S.3, when the sample size is small, neural networks tend to produce wider estimation intervals compared to random forests, reflecting greater estimation uncertainty. However, as both the sample size n and subsample size r increase, neural networks are able to capture the signal more effectively, leading to interval lengths and coverage that are comparable to, or even better than those of random forests. The

findings are consistent with the modern generalization theory of neural networks (Cao et al. 2022, Kou et al. 2023), which suggests that neural networks can generalize well when data quality is high and sufficient samples are available. In contrast, random forests tend to offer stronger robustness in low-data or high-noise scenarios.

To further elaborate on these points, we conduct additional experiments under a modified simulation setting, where we increase the signal strength and reduce the data noise. Specifically, we generate the covariates \mathbf{x}_i from a multivariate normal distribution $\mathcal{N}(0, \frac{1}{4}\mathbf{I}_{10})$, and define the underlying regression function as $g(\mathbf{x}) = 2x_1 + x_2^2 + 0.1 \cdot \arctan(x_3 - 0.3)$. The other settings remain unchanged. This setup introduces weaker noise component and stronger signal component.

The simulation results are summarized in Tables S.4 and S.5. In the logistic regression setting, neural networks tend to yield smaller estimation bias and narrower confidence intervals, while also achieving slightly higher coverage probabilities compared to random forests. This shows that, in scenarios with relatively strong signal structures and adequate subsample sizes, neural networks will offer improved estimation efficiency. For the Poisson model, both methods perform comparably in terms of coverage. However, as the subsample size r increases, the interval lengths produced by neural networks generally decrease, whereas those from random forests tend to increase. Overall, both methods have their strengths: RFs offer stability and robustness under noise, while DNNs can capture more complex patterns when data quality allows.

Finally, in all of these settings, we find that the corrected variance formula performs well for random forests. The corrected standard errors align more closely with the empirical standard deviations compared to the uncorrected ones. As a result, the coverage probabilities achieved by random forests are close to the nominal level.

A.3 Additional Experiments on Kernel Regressions

We consider a kernelized generalized linear model (GLM) to evaluate ESM under smooth nonlinear structures. Let $\gamma > 0$ and $k(\mathbf{x}, \mathbf{x}') = \exp\{-\gamma\|\mathbf{x} - \mathbf{x}'\|_2^2\}$ denote the Gaussian kernel. To make the kernel GLM computationally tractable for moderate to large sample sizes, we adopt the Nystrom approximation to obtain a finite dimensional feature representation $\phi(\mathbf{x}) \in \mathbb{R}^m$, where $k(\mathbf{x}, \mathbf{x}')$ can be approximated by $k(\mathbf{x}, \mathbf{x}') \approx \phi(\mathbf{x})^\top \phi(\mathbf{x}')$. Under this approximation, the kernel function becomes $f(\mathbf{x}) \approx \mathbf{w}^\top \phi(\mathbf{x})$, and the regularized empirical loss reduces to

$$\hat{\mathbf{w}} = \operatorname{argmin}_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n [\psi(\mathbf{w}^\top \phi(\mathbf{x}_i)) - y_i \mathbf{w}^\top \phi(\mathbf{x}_i)] + \lambda \|\mathbf{w}\|_2^2.$$

This linearized formulation retains the nonlinear representational capacity of the kernel model while allowing efficient estimation via standard GLM solvers. The conditional mean of the response can be written as $\mathbb{E}(y|\mathbf{x}) = \psi'(\hat{\mathbf{w}}^\top \phi(\mathbf{x}))$. In practice, we set $\gamma = 1/p$, where p is the input dimension to adapt the kernel smooth to feature scale. For the Bernoulli case, we fit a logistic regression with an L_2 penalty (inverse regularization strength $C = 10$) and the `lbfgs` solver with a maximum of 300 iterations. For the Poisson case, we employ the same Nystrom mapping followed by a Poisson regression with L_2 regularization parameter $\alpha = 10^{-3}$ and 300 iterations.

As shown in Table S.6, the kernel based models under the ESM framework also achieve empirical coverage probabilities close to the nominal 95% level. However, compared with the random forest and deep neural network (DNN) results, the kernel approach exhibits a higher MAE_f and wider average interval lengths, suggesting that the kernel regression is less effective in capturing complex nonlinear structures than the DNN.

Table S.6: Simulation Results with Kernel Regression.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.05(0.53)	0.43(0.32)	0.01(0.11)	0.09(0.07)	0.50	0.59	0.51	93.3%	0.41(0.07)
$r = n^{0.85}$	0.04(0.62)	0.50(0.37)	0.01(0.13)	0.10(0.08)	0.59	0.68	0.60	93.5%	0.45(0.09)
$r = n^{0.9}$	0.04(0.69)	0.56(0.42)	0.01(0.14)	0.12(0.08)	0.67	0.79	0.69	93.9%	0.50(0.11)
$n = 700, r = n^{0.8}$	0.04(0.50)	0.40(0.30)	0.01(0.11)	0.08(0.06)	0.47	0.59	0.47	93.0%	0.37(0.07)
$r = n^{0.85}$	0.02(0.57)	0.46(0.35)	0.00(0.12)	0.10(0.07)	0.54	0.69	0.55	93.3%	0.42(0.09)
$r = n^{0.9}$	0.01(0.63)	0.51(0.38)	0.00(0.13)	0.10(0.08)	0.60	0.78	0.62	93.8%	0.46(0.10)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.29(0.35)	0.37(0.27)	-0.18(0.25)	0.23(0.19)	0.32	0.39	0.33	84.7%	0.83(0.46)
$r = n^{0.85}$	-0.22(0.37)	0.34(0.26)	-0.13(0.27)	0.23(0.20)	0.35	0.41	0.36	89.5%	0.98(0.58)
$r = n^{0.9}$	-0.17(0.38)	0.33(0.25)	-0.09(0.30)	0.23(0.21)	0.36	0.43	0.38	91.6%	1.12(0.72)
$n = 700, r = n^{0.8}$	-0.20(0.28)	0.28(0.21)	-0.13(0.21)	0.19(0.16)	0.26	0.34	0.27	86.7%	0.71(0.37)
$r = n^{0.85}$	-0.14(0.30)	0.26(0.20)	-0.08(0.23)	0.18(0.16)	0.28	0.35	0.28	90.6%	0.81(0.45)
$r = n^{0.9}$	-0.10(0.30)	0.26(0.19)	-0.05(0.24)	0.18(0.16)	0.29	0.36	0.29	91.9%	0.87(0.50)

A.4 Additional Experiments on Dimensionality and Nonlinearity

We investigate how data dimensionality and nonlinearity influence the learning behavior of neural networks. Model performance depends not only on network capacity, but also on data quality, sample size, signal-to-noise ratio, and the complexity of the underlying regression function. We begin by assessing the effect of data dimensionality, as the inclusion of irrelevant covariates can increase noise and obscure the true signal. To quantify this effect, we compare three scenarios: low ($p = 7$), moderate ($p = 10$), and high ($p = 20$) dimensional settings. The corresponding results are summarized in Tables S.7 and S.8.

Table S.7: Simulation results with neural networks under $p = 20$.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.18(1.10)	0.90(0.66)	0.02(0.18)	0.15(0.10)	0.92	1.11	0.93	87.5%	0.54(0.20)
$r = n^{0.85}$	0.16(1.07)	0.86(0.65)	0.02(0.18)	0.15(0.10)	0.92	1.18	0.98	90.5%	0.58(0.20)
$r = n^{0.9}$	0.13(0.99)	0.79(0.61)	0.02(0.17)	0.14(0.10)	0.89	1.16	0.96	91.9%	0.58(0.19)
$n = 700, r = n^{0.8}$	0.13(0.83)	0.67(0.50)	0.02(0.14)	0.12(0.08)	0.69	0.95	0.71	88.4%	0.46(0.16)
$r = n^{0.85}$	0.08(0.69)	0.55(0.42)	0.01(0.13)	0.10(0.08)	0.62	0.86	0.64	92.1%	0.45(0.13)
$r = n^{0.9}$	0.05(0.57)	0.45(0.35)	0.01(0.11)	0.09(0.07)	0.52	0.73	0.54	93.3%	0.40(0.11)
$n = 2000, r = n^{0.8}$	0.08(0.54)	0.44(0.33)	0.01(0.10)	0.08(0.06)	0.47	0.86	0.48	90.9%	0.35(0.11)
$r = n^{0.85}$	0.03(0.49)	0.39(0.30)	0.00(0.10)	0.08(0.06)	0.46	0.85	0.47	93.0%	0.36(0.10)
$r = n^{0.9}$	0.02(0.40)	0.31(0.24)	0.00(0.08)	0.07(0.05)	0.38	0.72	0.40	95.2%	0.32(0.08)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.60(0.49)	0.63(0.45)	-0.30(0.29)	0.33(0.26)	0.43	0.53	0.44	77.1%	0.90(0.54)
$r = n^{0.85}$	-0.49(0.49)	0.55(0.42)	-0.25(0.31)	0.30(0.26)	0.44	0.58	0.48	86.6%	1.11(0.66)
$r = n^{0.9}$	-0.37(0.51)	0.49(0.40)	-0.17(0.35)	0.29(0.26)	0.48	0.62	0.51	90.7%	1.37(0.85)
$n = 700, r = n^{0.8}$	-0.46(0.39)	0.49(0.35)	-0.25(0.26)	0.28(0.23)	0.35	0.48	0.35	78.3%	0.79(0.47)
$r = n^{0.85}$	-0.31(0.38)	0.38(0.30)	-0.18(0.27)	0.24(0.22)	0.35	0.50	0.37	88.5%	0.95(0.56)
$r = n^{0.9}$	-0.20(0.38)	0.33(0.27)	-0.11(0.30)	0.23(0.22)	0.36	0.53	0.38	92.8%	1.11(0.68)
$n = 2000, r = n^{0.8}$	-0.22(0.26)	0.27(0.20)	-0.14(0.20)	0.19(0.17)	0.24	0.46	0.25	86.6%	0.67(0.39)
$r = n^{0.85}$	-0.15(0.27)	0.24(0.19)	-0.10(0.22)	0.17(0.17)	0.26	0.49	0.26	91.7%	0.78(0.46)
$r = n^{0.9}$	-0.08(0.26)	0.22(0.17)	-0.05(0.24)	0.17(0.17)	0.25	0.51	0.27	94.3%	0.85(0.52)

Table S.8: Simulation results with neural networks under $p = 7$.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.05(0.48)	0.38(0.30)	0.01(0.10)	0.08(0.06)	0.44	0.53	0.45	93.3%	0.34(0.09)
$r = n^{0.85}$	0.03(0.43)	0.34(0.27)	0.01(0.09)	0.07(0.06)	0.40	0.47	0.40	93.7%	0.31(0.08)
$r = n^{0.9}$	0.05(0.38)	0.30(0.24)	0.01(0.08)	0.06(0.05)	0.34	0.43	0.36	94.6%	0.29(0.07)
$n = 700, r = n^{0.8}$	0.04(0.30)	0.23(0.19)	0.01(0.06)	0.05(0.04)	0.27	0.36	0.28	93.9%	0.23(0.06)
$r = n^{0.85}$	0.03(0.30)	0.23(0.19)	0.01(0.06)	0.05(0.04)	0.27	0.38	0.28	93.5%	0.23(0.05)
$r = n^{0.9}$	0.04(0.31)	0.24(0.20)	0.01(0.06)	0.05(0.04)	0.26	0.35	0.27	92.2%	0.22(0.05)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.20(0.30)	0.28(0.23)	-0.12(0.22)	0.18(0.17)	0.27	0.34	0.28	90.6%	0.74(0.42)
$r = n^{0.85}$	-0.12(0.28)	0.23(0.20)	-0.07(0.22)	0.16(0.16)	0.26	0.33	0.27	93.7%	0.78(0.47)
$r = n^{0.9}$	-0.07(0.26)	0.21(0.17)	-0.04(0.22)	0.15(0.16)	0.25	0.31	0.26	94.6%	0.78(0.48)
$n = 700, r = n^{0.8}$	-0.09(0.20)	0.17(0.15)	-0.06(0.16)	0.12(0.12)	0.19	0.26	0.19	92.8%	0.55(0.32)
$r = n^{0.85}$	-0.03(0.19)	0.15(0.13)	-0.03(0.16)	0.11(0.11)	0.17	0.26	0.18	93.2%	0.52(0.26)
$r = n^{0.9}$	-0.02(0.20)	0.15(0.12)	-0.02(0.17)	0.12(0.12)	0.18	0.26	0.19	94.4%	0.59(0.34)

Table S.9: Simulation results with neural networks under different true function $g(\mathbf{x})$.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.04(0.79)	0.61(0.50)	0.01(0.15)	0.12(0.09)	0.72	0.90	0.75	93.5%	0.50(0.16)
$r = n^{0.85}$	0.06(0.70)	0.55(0.44)	0.01(0.14)	0.11(0.09)	0.65	0.78	0.66	93.6%	0.46(0.14)
$r = n^{0.9}$	0.07(0.60)	0.47(0.38)	0.01(0.12)	0.10(0.12)	0.54	0.68	0.57	93.8%	0.42(0.12)
$n = 700, r = n^{0.8}$	0.07(0.51)	0.40(0.32)	0.01(0.11)	0.08(0.07)	0.45	0.58	0.45	91.9%	0.34(0.10)
$r = n^{0.85}$	0.07(0.45)	0.35(0.28)	0.01(0.10)	0.07(0.06)	0.38	0.51	0.39	91.8%	0.31(0.09)
$r = n^{0.9}$	0.09(0.38)	0.31(0.24)	0.02(0.08)	0.07(0.05)	0.32	0.45	0.34	92.0%	0.28(0.08)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.42(0.50)	0.50(0.42)	-0.23(0.34)	0.29(0.29)	0.42	0.52	0.44	85.6%	0.97(0.66)
$r = n^{0.85}$	-0.29(0.48)	0.42(0.36)	-0.16(0.33)	0.26(0.26)	0.41	0.52	0.43	90.3%	1.09(0.76)
$r = n^{0.9}$	-0.18(0.45)	0.37(0.32)	-0.10(0.35)	0.25(0.27)	0.40	0.51	0.42	92.7%	1.22(0.93)
$n = 700, r = n^{0.8}$	-0.21(0.37)	0.32(0.28)	-0.14(0.27)	0.21(0.22)	0.30	0.41	0.31	88.0%	0.78(0.52)
$r = n^{0.85}$	-0.12(0.35)	0.28(0.24)	-0.08(0.28)	0.20(0.21)	0.29	0.40	0.30	90.5%	0.83(0.57)
$r = n^{0.9}$	-0.05(0.34)	0.26(0.22)	-0.05(0.27)	0.19(0.20)	0.27	0.39	0.28	91.2%	0.84(0.58)

As shown in Table S.7, increasing the dimension to $p = 20$ degrades estimation accuracy, with MAE_f rising from 0.28 to 0.55 and empirical standard deviation from 0.34 to 0.62 under the logistic regression model with $n = 700$. This reflects increased bias and variability, leading to a drop in coverage from 93.7% to 92.1%. Similar effects of high-dimensional noise on random forests were observed in Wager & Athey (2018). Increasing the sample size to $n = 2000$ under $p = 20$ improves estimation, restoring coverage to nearly 95%. Conversely, reducing the dimension to $p = 7$ improves neural network performance (Table S.8), with MAE_f decreasing to 0.23 and empirical standard deviation to 0.27. The reduced bias yields shorter intervals and coverage close to 95%.

We also consider another nonlinear function

$$g(\mathbf{x}) = 2 \tanh \left(\frac{1.5x_1 + 0.6(x_2^2 - 1) + 0.4x_3 \cdot \tanh(x_4) + 0.15 \cdot \sin(x_5)}{2.5} \right),$$

where $\tanh(x) = (e^x - e^{-x})/(e^x + e^{-x})$. This function introduces quadratic, tanh, and even interaction components, representing a moderately nonlinear structure that remains learnable for neural networks. All other simulation settings are kept the same as the basic

settings as those in Section 4.

Table S.9 shows the results with a different function $g(\mathbf{x})$. We find that the corrected standard error estimator remains accurate, but the coverage probability shows a mild drop compared to the basic setting. This phenomenon shows that not only random forests but also neural networks may exhibit bias in their output logits, particularly under complex nonlinear structures. Neural networks, especially in nonlinear and high-dimensional regimes, can incur non-negligible approximation bias due to optimization challenges or limitations in representing the true underlying function $g(\mathbf{x})$. Such biases are difficult to eliminate in practice and can result in reduced coverage probability, even when variance estimation is accurate. How to effectively reduce or correct for such bias remains an open problem, and addressing it is key to improving the reliability of inference.

A.5 Additional Experiments with Optimization

We examine how optimization hyperparameters, i.e., weight decay and training epochs, affect results and find our inference stable across reasonable settings. In particular, we vary training epochs (300, 700) and weight decay (2×10^{-3} , 2×10^{-4}).

As shown in Tables S.10 and S.11, our variance estimator remains accurate across different training epochs and weight decay values, closely matching the empirical standard deviation. This confirms that optimization error is negligible under reasonable hyperparameter settings and supports viewing Assumption 3.3 as a technical device separating optimization from inference rather than a practical constraint on SGD convergence. Larger subsample sizes r reduce bias in MAE, consistent with modern generalization theory. Overall, the results validate the robustness of our inference guarantees, while highlighting that extreme hyperparameters (e.g., excessive weight decay or too few epochs) can impair learning and

Table S.10: Simulation results with different epochs.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic with Epochs = 300									
$n = 400, r = n^{0.8}$	0.08(0.50)	0.40(0.31)	0.01(0.10)	0.08(0.06)	0.46	0.56	0.47	93.4%	0.36(0.09)
$r = n^{0.85}$	0.07(0.44)	0.35(0.28)	0.02(0.09)	0.07(0.06)	0.41	0.50	0.42	93.9%	0.34(0.08)
$r = n^{0.9}$	0.07(0.40)	0.32(0.26)	0.01(0.09)	0.07(0.06)	0.38	0.46	0.39	93.9%	0.32(0.07)
$n = 700, r = n^{0.8}$	0.07(0.33)	0.26(0.21)	0.01(0.07)	0.06(0.05)	0.29	0.39	0.30	92.5%	0.25(0.05)
$r = n^{0.85}$	0.07(0.31)	0.25(0.19)	0.02(0.07)	0.05(0.04)	0.27	0.36	0.27	91.6%	0.23(0.05)
$r = n^{0.9}$	0.06(0.30)	0.24(0.19)	0.01(0.07)	0.05(0.04)	0.25	0.35	0.26	90.0%	0.22(0.05)
Logistic with Epochs = 700									
$n = 400, r = n^{0.8}$	0.13(0.86)	0.69(0.53)	0.02(0.15)	0.13(0.09)	0.76	0.93	0.78	90.5%	0.51(0.15)
$r = n^{0.85}$	0.08(0.72)	0.57(0.45)	0.01(0.14)	0.11(0.08)	0.67	0.83	0.70	92.8%	0.49(0.14)
$r = n^{0.9}$	0.06(0.61)	0.48(0.38)	0.01(0.12)	0.10(0.07)	0.58	0.72	0.60	94.1%	0.45(0.12)
$n = 700, r = n^{0.8}$	0.07(0.50)	0.40(0.31)	0.01(0.10)	0.08(0.06)	0.46	0.62	0.47	92.9%	0.36(0.09)
$r = n^{0.85}$	0.05(0.41)	0.32(0.25)	0.01(0.09)	0.07(0.05)	0.39	0.52	0.40	94.5%	0.32(0.08)
$r = n^{0.9}$	0.05(0.36)	0.28(0.22)	0.01(0.08)	0.06(0.05)	0.34	0.46	0.35	94.5%	0.29(0.07)
Poisson with Epochs = 300									
$n = 400, r = n^{0.8}$	-0.23(0.32)	0.30(0.25)	-0.13(0.22)	0.19(0.16)	0.30	0.37	0.31	90.2%	0.78(0.40)
$r = n^{0.85}$	-0.15(0.31)	0.27(0.22)	-0.08(0.23)	0.18(0.16)	0.30	0.37	0.30	92.7%	0.84(0.45)
$r = n^{0.9}$	-0.08(0.30)	0.24(0.20)	-0.04(0.23)	0.17(0.16)	0.29	0.36	0.30	94.2%	0.88(0.48)
$n = 700, r = n^{0.8}$	-0.12(0.23)	0.20(0.16)	-0.07(0.17)	0.14(0.13)	0.21	0.29	0.22	92.2%	0.60(0.30)
$r = n^{0.85}$	-0.06(0.22)	0.18(0.14)	-0.04(0.17)	0.13(0.12)	0.21	0.29	0.21	93.8%	0.62(0.31)
$r = n^{0.9}$	-0.01(0.22)	0.17(0.14)	-0.01(0.17)	0.13(0.12)	0.20	0.28	0.20	93.7%	0.62(0.31)
Poisson with Epochs = 700									
$n = 400, r = n^{0.8}$	-0.37(0.41)	0.43(0.34)	-0.19(0.24)	0.24(0.19)	0.37	0.47	0.39	87.2%	0.91(0.49)
$r = n^{0.85}$	-0.26(0.39)	0.36(0.30)	-0.13(0.26)	0.22(0.19)	0.37	0.47	0.39	92.0%	1.02(0.57)
$r = n^{0.9}$	-0.18(0.37)	0.32(0.26)	-0.08(0.26)	0.21(0.18)	0.35	0.46	0.38	93.9%	1.10(0.65)
$n = 700, r = n^{0.8}$	-0.20(0.29)	0.27(0.22)	-0.11(0.20)	0.17(0.15)	0.27	0.37	0.28	90.3%	0.72(0.37)
$r = n^{0.85}$	-0.12(0.27)	0.23(0.19)	-0.07(0.20)	0.16(0.14)	0.26	0.36	0.27	93.1%	0.76(0.40)
$r = n^{0.9}$	-0.06(0.26)	0.21(0.17)	-0.03(0.21)	0.15(0.14)	0.25	0.35	0.26	94.3%	0.77(0.42)

should be avoided.

A.6 Additional Experiments on r

Recall that $r = n^\gamma$. We vary γ from 0.65 to 0.95 (step 0.05) while keeping other settings as in Section 4, and plot the standard deviation, MAE of $\widehat{f}^B - f_0$, and coverage probability versus γ for $n = 400$ and $n = 700$ under two models.

Figures S.2a–S.2d show that as γ increases, our IJ variance estimator closely matches the empirical standard deviation in both logistic and Poisson models, confirming its robust-

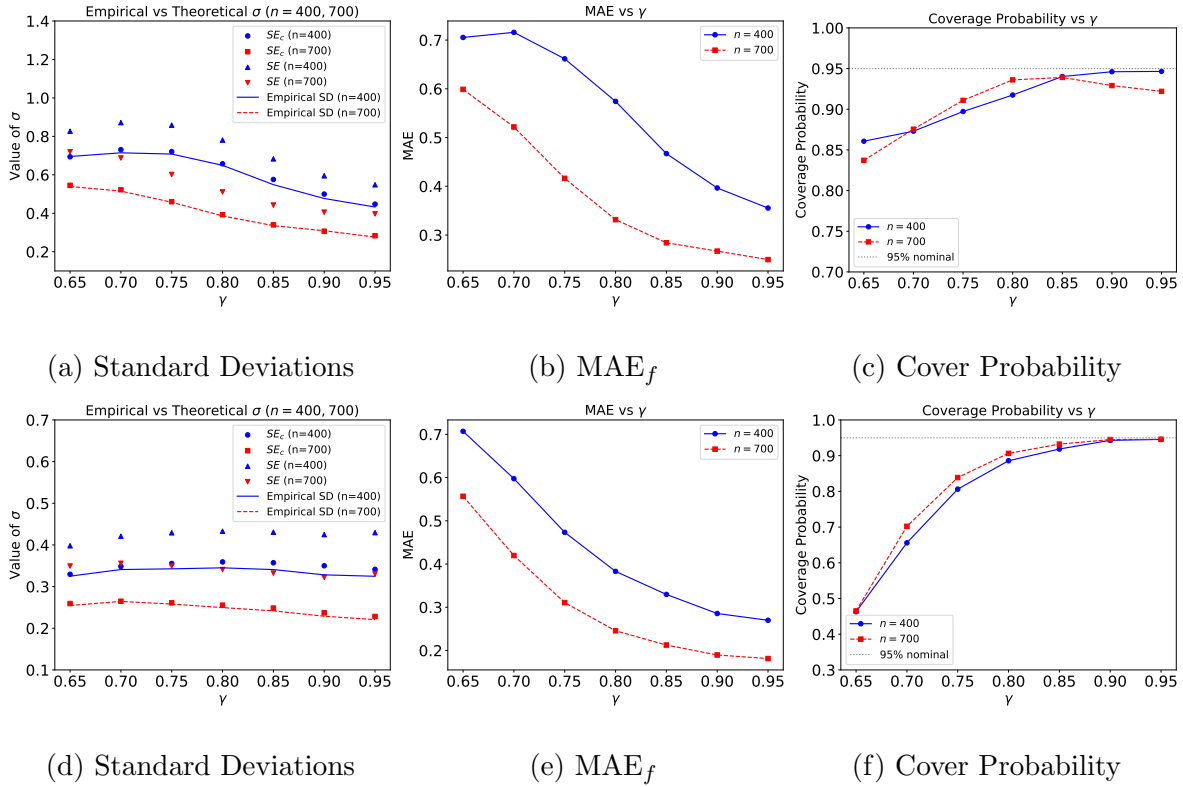


Figure S.2: Simulation results with different values of γ . Blue and red lines/points correspond to $n = 400$ and $n = 700$ respectively. Figures S.2a–S.2c present the results under the logistic model, and Figures S.2d–S.2f show results under the Poisson model.

Table S.11: Simulation results with different weight decay.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic with WeightDecay = 2×10^{-3}									
$n = 400, r = n^{0.8}$	0.18(1.61)	1.28(0.98)	0.02(0.22)	0.19(0.11)	1.25	1.54	1.27	87.5%	0.62(0.22)
$r = n^{0.85}$	0.13(1.30)	1.02(0.81)	0.02(0.20)	0.16(0.10)	1.08	1.36	1.13	90.9%	0.61(0.19)
$r = n^{0.9}$	0.08(1.06)	0.83(0.66)	0.01(0.18)	0.15(0.10)	0.94	1.19	0.97	92.9%	0.59(0.17)
$n = 700, r = n^{0.8}$	0.10(0.93)	0.73(0.58)	0.01(0.15)	0.13(0.08)	0.75	1.02	0.76	89.2%	0.48(0.15)
$r = n^{0.85}$	0.07(0.72)	0.57(0.45)	0.01(0.13)	0.11(0.08)	0.63	0.87	0.64	92.0%	0.45(0.12)
$r = n^{0.9}$	0.03(0.59)	0.46(0.37)	0.00(0.11)	0.09(0.07)	0.54	0.76	0.55	93.8%	0.41(0.10)
Logistic with WeightDecay = 2×10^{-4}									
$n = 400, r = n^{0.8}$	0.22(1.72)	1.36(1.07)	0.02(0.22)	0.19(0.11)	1.32	1.66	1.37	88.1%	0.64(0.23)
$r = n^{0.85}$	0.15(1.41)	1.11(0.88)	0.02(0.20)	0.17(0.11)	1.17	1.48	1.22	91.1%	0.64(0.20)
$r = n^{0.9}$	0.09(1.13)	0.88(0.71)	0.01(0.18)	0.15(0.10)	1.00	1.28	1.04	93.3%	0.61(0.17)
$n = 700, r = n^{0.8}$	0.12(1.00)	0.80(0.63)	0.02(0.16)	0.14(0.09)	0.80	1.10	0.82	89.1%	0.50(0.16)
$r = n^{0.85}$	0.07(0.78)	0.61(0.49)	0.01(0.14)	0.11(0.08)	0.67	0.93	0.69	92.0%	0.47(0.13)
$r = n^{0.9}$	0.04(0.64)	0.50(0.40)	0.01(0.12)	0.10(0.07)	0.59	0.83	0.60	93.6%	0.43(0.11)
Poisson with WeightDecay = 2×10^{-3}									
$n = 400, r = n^{0.8}$	-0.67(0.60)	0.70(0.56)	-0.28(0.25)	0.31(0.21)	0.49	0.63	0.51	83.0%	0.95(0.51)
$r = n^{0.85}$	-0.49(0.56)	0.56(0.49)	-0.21(0.27)	0.27(0.20)	0.48	0.62	0.50	88.8%	1.10(0.60)
$r = n^{0.9}$	-0.35(0.50)	0.46(0.41)	-0.14(0.29)	0.25(0.20)	0.46	0.61	0.49	92.9%	1.24(0.72)
$n = 700, r = n^{0.8}$	-0.41(0.41)	0.44(0.37)	-0.20(0.21)	0.23(0.17)	0.35	0.49	0.36	84.3%	0.77(0.39)
$r = n^{0.85}$	-0.27(0.37)	0.35(0.30)	-0.13(0.22)	0.20(0.16)	0.33	0.48	0.34	90.5%	0.84(0.43)
$r = n^{0.9}$	-0.18(0.35)	0.30(0.26)	-0.08(0.24)	0.19(0.16)	0.33	0.48	0.33	92.8%	0.90(0.49)
Poisson with WeightDecay = 2×10^{-4}									
$n = 400, r = n^{0.8}$	-0.72(0.64)	0.75(0.61)	-0.29(0.25)	0.32(0.22)	0.52	0.65	0.53	81.6%	0.96(0.51)
$r = n^{0.85}$	-0.54(0.59)	0.60(0.53)	-0.22(0.27)	0.29(0.21)	0.51	0.65	0.52	88.4%	1.11(0.61)
$r = n^{0.9}$	-0.38(0.52)	0.48(0.42)	-0.16(0.29)	0.26(0.21)	0.47	0.64	0.51	92.7%	1.26(0.73)
$n = 700, r = n^{0.8}$	-0.44(0.43)	0.48(0.39)	-0.21(0.21)	0.24(0.17)	0.36	0.51	0.37	83.1%	0.77(0.39)
$r = n^{0.85}$	-0.30(0.39)	0.38(0.33)	-0.14(0.23)	0.21(0.17)	0.35	0.50	0.36	89.6%	0.85(0.44)
$r = n^{0.9}$	-0.20(0.36)	0.31(0.27)	-0.09(0.24)	0.19(0.16)	0.33	0.50	0.34	92.9%	0.92(0.50)

ness and validating the finite-sample correction factor $n(n-1)/(n-r)^2$. Figures S.2b–S.2e indicate that MAE decreases with larger r , supporting Theorem 3.5 and showing bias reduction with larger subsamples. Finally, Figures S.2c–S.2f demonstrate that as r increases, coverage improves toward the nominal 95% level, consistent with our assumption on γ ensuring bias decay and valid inference.

We also evaluate computational cost across different r values by recording the training time of a single neural network. All the simulations are conducted in the institute high performance computing platform. The simulation is repeated 300 times, with B networks

Table S.12: Computation time in seconds.

$\gamma =$	0.65	0.7	0.75	0.8	0.85	0.9	0.95
Logistic	0.74	0.82	1.00	1.38	1.41	1.90	2.30
Poisson	0.75	0.90	0.96	1.32	1.41	1.78	2.24


trained per repetition. As shown in Table S.12, training one network takes only a few seconds, while the main cost arises from repeatedly training the B networks. Following Wager & Athey (2018) and Fei & Li (2024), we let B grow proportionally with n and apply bias correction to maintain variance estimation accuracy, largely reducing computation. Although $300 \times B$ trainings would otherwise be intensive, high-performance computing with parallel nodes and multi-CPU systems makes the experiments feasible.

A.7 Additional Experiments on B

We decrease B from 1400 to 400 in increments of 200 to assess the consistency of variance estimation, while fixing $\gamma = 0.9$.

Figure S.3 reports the standard deviations, MAE_f , and coverage probabilities as B varies, analogous to Section A.6. As B increases, the gap between empSD and the uncorrected SE narrows, while our bias-corrected estimator SE_c closely matches empSD even when $B = 400$ and $n = 700$ ($B < n$). The MAE_f is smaller for $n = 700$ than for $n = 400$, and all coverage probabilities remain near the nominal 95% level. These results show that the bias correction is effective even with smaller B , allowing valid inference and stable performance while reducing computational cost.


We further assess the impact of a larger subsample size ($B = 3000$) on variance estimation and the accuracy of the correction term. Table S.13 shows that the results remain consistent: increasing B reduces MAE for both \hat{f}^B and $\psi'(\hat{f}^B)$, and the empirical standard deviations (empSD) align closely with the estimated standard errors (SE), maintaining



(a) Standard Deviations

(b) MAE_f

(c) Cover Probability



(d) Standard Deviations

(e) MAE_f

(f) Cover Probability

Figure S.3: Simulation results with different values of B . Blue and red lines/points correspond to $n = 400$ and $n = 700$ respectively. Figures S.3a–S.3c present the results under the logistic model, and Figures S.3d–S.3f show results under the Poisson model.

Table S.13: Simulation results for different models under varying sample sizes n and subsampling ratios r with larger $B = 3000$.


	Bias _{f}	MAE _{f}	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _{c}	CP	AIL
Logistic Model									
$n = 400, r = n^{0.8}$	0.10(0.71)	0.56(0.44)	0.02(0.13)	0.11(0.08)	0.64	0.72	0.66	92.0%	0.46(0.13)
$r = n^{0.85}$	0.07(0.60)	0.47(0.37)	0.01(0.12)	0.10(0.07)	0.56	0.63	0.57	93.4%	0.43(0.11)
$r = n^{0.9}$	0.05(0.49)	0.39(0.31)	0.01(0.10)	0.08(0.06)	0.47	0.55	0.50	95.3%	0.39(0.09)
$n = 700, r = n^{0.8}$	0.06(0.41)	0.33(0.26)	0.01(0.09)	0.07(0.05)	0.38	0.46	0.40	93.7%	0.31(0.08)
$r = n^{0.85}$	0.05(0.35)	0.28(0.22)	0.01(0.08)	0.06(0.05)	0.33	0.39	0.34	94.2%	0.28(0.06)
$r = n^{0.9}$	0.05(0.33)	0.26(0.20)	0.01(0.07)	0.06(0.04)	0.30	0.36	0.31	93.8%	0.26(0.06)
Poisson Model									
$n = 400, r = n^{0.8}$	-0.32(0.38)	0.39(0.32)	-0.17(0.24)	0.22(0.18)	0.36	0.39	0.36	88.1%	0.87(0.46)
$r = n^{0.85}$	-0.21(0.37)	0.32(0.27)	-0.10(0.25)	0.20(0.18)	0.34	0.39	0.36	92.2%	0.96(0.53)
$r = n^{0.9}$	-0.13(0.35)	0.28(0.24)	-0.06(0.26)	0.20(0.18)	0.33	0.38	0.35	93.5%	1.02(0.58)
$n = 700, r = n^{0.8}$	-0.16(0.27)	0.24(0.20)	-0.09(0.19)	0.16(0.14)	0.25	0.30	0.25	90.4%	0.68(0.34)
$r = n^{0.85}$	-0.10(0.26)	0.21(0.17)	-0.06(0.19)	0.15(0.14)	0.24	0.30	0.25	93.1%	0.71(0.37)
$r = n^{0.9}$	-0.04(0.24)	0.19(0.15)	-0.02(0.19)	0.14(0.13)	0.23	0.28	0.24	95.2%	0.72(0.38)

nominal 95% coverage. Figures S.4a and S.4b illustrate that our estimator accurately recovers the conditional mean and that the corrected variance closely follows empSD, whereas the uncorrected version slightly overestimates. As B increases from 1400 to 3000, this bias diminishes, consistent with theory showing that the correction term vanishes asymptotically. These findings confirm that the correction improves finite-sample accuracy while preserving asymptotic validity.

A.8 Additional Experiments on L

We now examine how network depth affects estimation bias and variance using $n = 700$, varying both L and r . For each L , we set $\mathbf{p} = (p, \underbrace{30, \dots, 30}_{L-1}, 1)$ and apply batch normalization after the 2nd, 4th, and 6th layers to ensure gradient stability. We consider $L = 3, 4, 5, 6$ and $r = n^{0.8}, n^{0.85}, n^{0.9}, n^{0.95}$.

Tables S.14–S.15 show that deeper networks tend to introduce larger biases, consistent with (C.2), which predicts larger approximation error and lower coverage for deeper archi-



(a) Point Estimates and Variations (b) Estimated Standard Errors and Variations

Figure S.4: Estimation and inference in simulation samples: Logistic Model with $n = 700$, $r = n^{0.9}$ and $B = 3000$. Figure S.4a shows the average estimated $\mathbb{E}(y|\mathbf{x})$ with variability across 300 repetitions (gray band) over test points. Figure S.4b compares corrected and uncorrected standard errors and their variability (gray and blue bands, respectively) to the empirical standard deviations of estimates across all test samples.

tectures. Hence, in our settings, it appears that relatively shallow networks are preferable for both stability and reduced bias (Goodfellow et al. 2016).

Consistent with the sensitivity analysis in Section A.6, increasing r reduces MAE_f , and all bias-corrected variance estimators closely match the empirical results, confirming the robustness of our inference procedure.

Table S.14: Simulation results with Logistic under varying L and r .

	Bias _{f}	MAE _{f}	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _{c}	CP	AIL
Logistic									
$r = n^{0.8}, L = 3$	0.19(0.66)	0.57(0.39)	0.04(0.13)	0.11(0.08)	0.59	0.92	0.60	84.5%	0.43(0.15)
$L = 4$	0.21(0.71)	0.61(0.41)	0.04(0.14)	0.12(0.08)	0.61	0.98	0.61	82.7%	0.44(0.16)
$L = 5$	0.21(0.73)	0.63(0.43)	0.04(0.14)	0.12(0.08)	0.63	1.11	0.63	82.0%	0.44(0.17)
$L = 6$	0.21(0.70)	0.61(0.41)	0.04(0.14)	0.12(0.08)	0.61	1.10	0.63	83.4%	0.44(0.16)
$r = n^{0.85}, L = 3$	0.19(0.69)	0.59(0.41)	0.04(0.14)	0.12(0.08)	0.63	1.03	0.64	85.6%	0.46(0.16)
$L = 4$	0.21(0.74)	0.64(0.43)	0.04(0.14)	0.13(0.08)	0.66	1.12	0.67	83.2%	0.46(0.17)
$L = 5$	0.20(0.80)	0.69(0.46)	0.04(0.15)	0.13(0.08)	0.71	1.32	0.71	81.8%	0.48(0.19)
$L = 6$	0.21(0.77)	0.66(0.45)	0.04(0.15)	0.13(0.09)	0.69	1.31	0.70	83.0%	0.48(0.18)
$r = n^{0.9}, L = 3$	0.19(0.71)	0.60(0.42)	0.04(0.14)	0.12(0.08)	0.65	1.16	0.68	86.8%	0.48(0.16)
$L = 4$	0.21(0.75)	0.64(0.44)	0.04(0.15)	0.13(0.08)	0.69	1.27	0.71	84.9%	0.49(0.18)
$L = 5$	0.21(0.85)	0.73(0.48)	0.04(0.16)	0.14(0.09)	0.76	1.60	0.78	81.9%	0.51(0.21)
$L = 6$	0.20(0.82)	0.72(0.47)	0.04(0.16)	0.14(0.09)	0.75	1.57	0.76	82.7%	0.51(0.20)
$r = n^{0.95}, L = 3$	0.17(0.71)	0.59(0.42)	0.03(0.14)	0.12(0.08)	0.66	1.37	0.69	88.9%	0.49(0.16)
$L = 4$	0.20(0.74)	0.63(0.44)	0.04(0.15)	0.12(0.09)	0.69	1.48	0.72	86.5%	0.49(0.18)
$L = 5$	0.19(0.90)	0.77(0.50)	0.03(0.17)	0.15(0.09)	0.82	2.14	0.84	81.7%	0.54(0.23)
$L = 6$	0.21(0.86)	0.73(0.49)	0.04(0.17)	0.14(0.09)	0.79	2.07	0.82	81.9%	0.53(0.23)

Table S.15: Simulation results with Poisson under varying L and r .

	Bias _{f}	MAE _{f}	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _{c}	CP	AIL
Poisson									
$r = n^{0.8}, L = 3$	-0.38(0.33)	0.41(0.29)	-0.20(0.21)	0.23(0.17)	0.31	0.49	0.31	79.8%	0.71(0.37)
$L = 4$	-0.42(0.37)	0.45(0.32)	-0.20(0.21)	0.24(0.17)	0.33	0.53	0.33	76.9%	0.74(0.40)
$L = 5$	-0.47(0.38)	0.50(0.34)	-0.23(0.22)	0.26(0.17)	0.34	0.59	0.34	72.9%	0.75(0.44)
$L = 6$	-0.48(0.39)	0.51(0.35)	-0.22(0.22)	0.26(0.17)	0.34	0.62	0.35	73.7%	0.78(0.46)
$r = n^{0.85}, L = 3$	-0.30(0.35)	0.36(0.28)	-0.15(0.22)	0.21(0.17)	0.33	0.55	0.33	87.0%	0.82(0.43)
$L = 4$	-0.35(0.38)	0.41(0.31)	-0.16(0.23)	0.22(0.17)	0.35	0.61	0.36	84.6%	0.86(0.46)
$L = 5$	-0.42(0.41)	0.48(0.34)	-0.19(0.24)	0.25(0.18)	0.37	0.69	0.37	79.6%	0.88(0.52)
$L = 6$	-0.43(0.42)	0.48(0.35)	-0.19(0.24)	0.25(0.18)	0.38	0.72	0.38	80.0%	0.90(0.53)
$r = n^{0.9}, L = 3$	-0.22(0.36)	0.33(0.27)	-0.10(0.24)	0.20(0.17)	0.34	0.63	0.35	91.2%	0.94(0.50)
$L = 4$	-0.26(0.39)	0.37(0.30)	-0.11(0.25)	0.21(0.17)	0.37	0.69	0.38	89.1%	0.98(0.52)
$L = 5$	-0.35(0.43)	0.44(0.34)	-0.15(0.26)	0.24(0.18)	0.40	0.82	0.40	84.8%	1.02(0.60)
$L = 6$	-0.36(0.45)	0.45(0.35)	-0.14(0.27)	0.25(0.18)	0.41	0.86	0.41	84.5%	1.05(0.62)
$r = n^{0.95}, L = 3$	-0.14(0.37)	0.30(0.25)	-0.04(0.27)	0.20(0.18)	0.35	0.77	0.37	93.1%	1.06(0.62)
$L = 4$	-0.17(0.40)	0.33(0.28)	-0.05(0.27)	0.21(0.18)	0.37	0.84	0.39	91.6%	1.11(0.62)
$L = 5$	-0.26(0.45)	0.41(0.32)	-0.09(0.29)	0.24(0.19)	0.42	1.06	0.43	87.8%	1.20(0.74)
$L = 6$	-0.27(0.46)	0.42(0.33)	-0.09(0.29)	0.24(0.19)	0.42	1.11	0.44	87.2%	1.22(0.76)

Table S.16: Simulation results with different dropout rate.

	Bias _f	MAE _f	Bias _{ψ'}	MAE _{ψ'}	EmpSD	SE	SE _c	CP	AIL
Logistic with Dropout Rate = 0.3									
$n = 400, r = n^{0.8}$	0.09(0.59)	0.47(0.37)	0.02(0.12)	0.10(0.07)	0.55	0.68	0.56	92.9%	0.41(0.11)
$r = n^{0.85}$	0.08(0.51)	0.41(0.32)	0.02(0.11)	0.09(0.07)	0.48	0.60	0.49	93.9%	0.38(0.09)
$r = n^{0.9}$	0.06(0.44)	0.35(0.28)	0.01(0.10)	0.08(0.06)	0.42	0.54	0.43	94.1%	0.35(0.08)
$n = 700, r = n^{0.8}$	0.08(0.35)	0.28(0.23)	0.02(0.08)	0.06(0.05)	0.33	0.46	0.33	93.2%	0.27(0.06)
$r = n^{0.85}$	0.06(0.32)	0.25(0.20)	0.01(0.04)	0.06(0.04)	0.29	0.41	0.30	93.1%	0.25(0.05)
$r = n^{0.9}$	0.07(0.31)	0.25(0.19)	0.02(0.07)	0.06(0.04)	0.27	0.39	0.27	90.9%	0.23(0.05)
Logistic with Dropout Rate = 0.5									
$n = 400, r = n^{0.8}$	0.11(0.48)	0.38(0.31)	0.02(0.10)	0.08(0.06)	0.45	0.58	0.46	93.8%	0.36(0.09)
$r = n^{0.85}$	0.08(0.42)	0.33(0.27)	0.02(0.09)	0.07(0.06)	0.40	0.52	0.40	93.9%	0.33(0.07)
$r = n^{0.9}$	0.07(0.39)	0.31(0.25)	0.02(0.09)	0.07(0.05)	0.36	0.48	0.36	93.0%	0.30(0.07)
$n = 700, r = n^{0.8}$	0.07(0.31)	0.25(0.20)	0.02(0.07)	0.06(0.04)	0.28	0.40	0.28	92.1%	0.24(0.05)
$r = n^{0.85}$	0.07(0.30)	0.24(0.19)	0.02(0.07)	0.05(0.04)	0.25	0.38	0.26	89.6%	0.22(0.05)
$r = n^{0.9}$	0.07(0.31)	0.25(0.19)	0.02(0.07)	0.06(0.04)	0.24	0.38	0.24	86.1%	0.21(0.04)
Poisson with Dropout Rate = 0.3									
$n = 400, r = n^{0.8}$	-0.24(0.33)	0.31(0.26)	-0.13(0.22)	0.19(0.16)	0.31	0.39	0.31	90.8%	0.77(0.36)
$r = n^{0.85}$	-0.13(0.31)	0.25(0.22)	-0.07(0.22)	0.17(0.15)	0.28	0.38	0.30	94.0%	0.82(0.40)
$r = n^{0.9}$	-0.07(0.29)	0.23(0.20)	-0.04(0.22)	0.16(0.15)	0.28	0.37	0.29	94.6%	0.85(0.43)
$n = 700, r = n^{0.8}$	-0.10(0.23)	0.19(0.16)	-0.06(0.17)	0.13(0.12)	0.21	0.30	0.22	92.8%	0.59(0.27)
$r = n^{0.85}$	-0.05(0.22)	0.17(0.14)	-0.03(0.17)	0.13(0.12)	0.20	0.30	0.21	93.6%	0.60(0.28)
$r = n^{0.9}$	-0.01(0.21)	0.16(0.13)	-0.01(0.17)	0.12(0.11)	0.19	0.30	0.20	93.6%	0.60(0.29)
Poisson with Dropout Rate = 0.5									
$n = 400, r = n^{0.8}$	-0.15(0.29)	0.25(0.22)	-0.09(0.20)	0.16(0.14)	0.26	0.35	0.27	92.9%	0.70(0.29)
$r = n^{0.85}$	-0.07(0.27)	0.21(0.18)	-0.04(0.20)	0.15(0.14)	0.25	0.33	0.25	93.8%	0.72(0.31)
$r = n^{0.9}$	-0.03(0.26)	0.21(0.17)	-0.02(0.20)	0.15(0.13)	0.24	0.34	0.25	93.6%	0.72(0.33)
$n = 700, r = n^{0.8}$	-0.05(0.21)	0.17(0.14)	-0.03(0.16)	0.12(0.11)	0.18	0.27	0.18	91.8%	0.51(0.21)
$r = n^{0.85}$	-0.01(0.20)	0.16(0.13)	-0.02(0.16)	0.11(0.11)	0.18	0.27	0.18	91.6%	0.51(0.22)
$r = n^{0.9}$	0.02(0.20)	0.16(0.12)	0.01(0.15)	0.11(0.10)	0.17	0.28	0.17	90.3%	0.51(0.22)

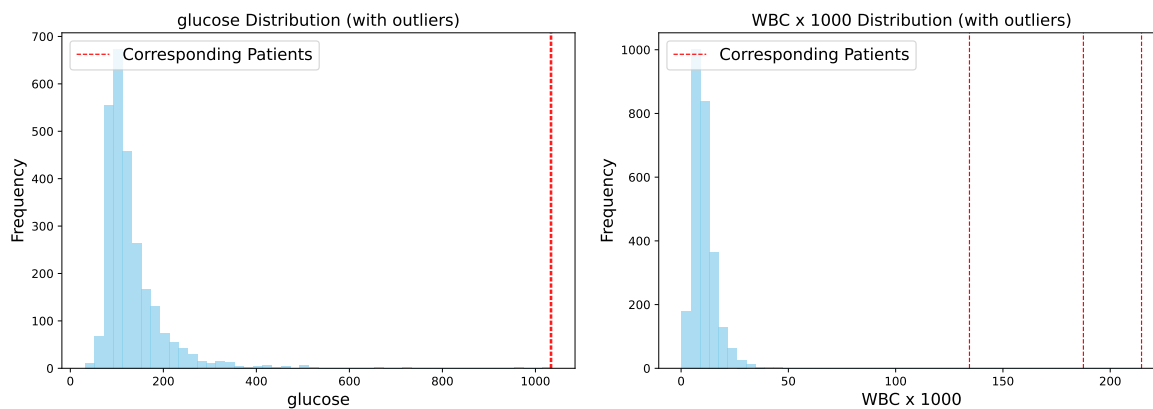
A.9 Additional Experiments on Dropout Rate

We vary the dropout rates, which are set to 0.3 and 0.5 in the two models, while retaining all other settings as in Section 4. As shown in Table S.16, under different dropout rates, both the bias and MAE_f decrease as r increases, showing a consistent trend across settings. The proposed variance estimator also remains accurate under all configurations. The coverage probability can be affected when the dropout rate is large, as excessive dropout may limit the network’s effective capacity and lead to mild underfitting.

B Additional Real Data Experiments

B.1 Additional Investigation on Large Volatility in DNNs

We further analyze patients with unusually wide prediction intervals from neural networks. As shown in Figure 4b, patients indexed around 800–1600 exhibit large intervals despite similar predicted risks. Among the five with the widest intervals, outlier values in glucose and WBC likely drive the elevated uncertainty. Figure S.5 highlights these extreme values.



(a) Histogram for “glucose”

(b) Histogram for “WBC”

Figure S.5: Distributions of glucose and WBC levels for the five distinct patients with the widest prediction intervals. Red dashed lines mark their observed values: two patients have extreme glucose values and three others have extreme WBC values, corresponding to those exhibiting the highest uncertainty under neural network estimations.

Figure S.5 shows that two patients have extremely high glucose levels and three have unusually high WBC, all well outside the typical range. For neural networks, such outliers reduce prediction confidence, leading to wider intervals, an adaptive and desirable feature for guarding against overconfidence in poorly supported regions. In contrast, random forests, while more robust to outliers due to localized splitting, may produce overly narrow

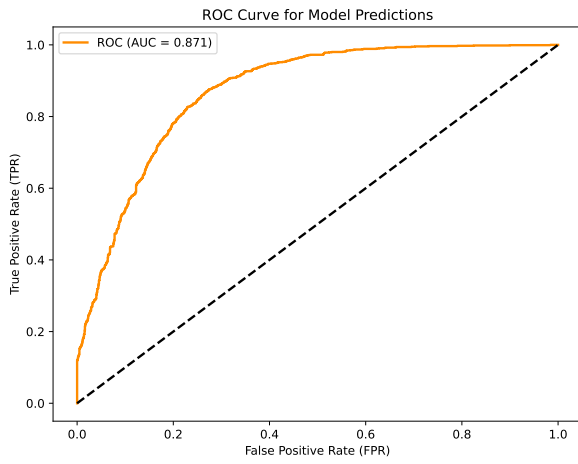
intervals, underestimating uncertainty in critical cases.

B.2 Additional Experiments with Additional Features


As a sensitivity analysis, we extend the real data experiments by adding features such as RDW, MCHC, MCV, and admission indicators (e.g., Respiratory, Sepsis-related, Surgical/Trauma), expanding the covariate set to 46. This evaluates how model performance and uncertainty estimates respond to a richer feature space with added clinical signal and potential noise. Figures S.6 and S.7 show results under the expanded feature set. For logistic regression, AUCs for DNNs and RFs remain similar, but DNNs adaptively widen their intervals (0.306 vs. 0.161), reflecting greater sensitivity to increased feature complexity. In the Poisson model, AUCs and interval lengths are comparable (0.261 for DNNs, 0.274 for RFs). These findings highlight the strength of DNNs in leveraging richer features and adapting uncertainty when signal dominates, while RFs maintain stability in noisier, high-dimensional settings, consistent with trends observed in simulation

B.3 Additional Experiments on Cross-Site Model Transferability


We further transferred the model trained from hospital site 188 to hospital site 458. Site 458 was selected because it has a comparable patient volume to the original training site (188) but exhibits a markedly different outcome distribution. Although the cohort sizes are similar, the risk profiles differ. As shown in Figure S.8, patients at site 458 generally tend to have a lower anticipated mortality likelihood (assessed by physicians at baseline), a main risk factor for ICU readmissions. Consequently, site 458 displays a lower overall readmission rate. Specifically, it includes 2,178 patients without readmission and 111 patients with one or more readmissions (ranging from 1 to 4 events), resulting in a challenging and imbalanced




(a) ROC curve in DNNs



(b) Individual estimates and CIs in DNNs




(c) ROC curve in RFs




(d) Individual estimates and CIs in RFs

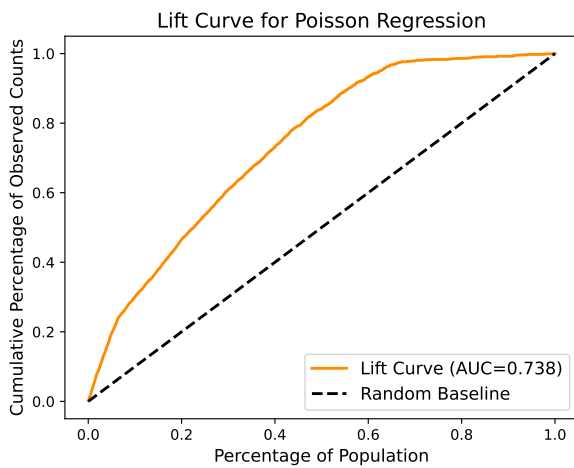
Figure S.6: Evaluation of the nonparametric logistic model estimator. Figure S.6a and S.6c show the ROC curve with both AUC of 0.84. Figure S.6b and S.6d display estimated subject-level probabilities of ICU readmission and confidence intervals, illustrating heteroskedasticity across individuals.




(a) Lift curve in DNNs



(b) Individual estimates and CIs in DNNs



(c) Lift curve in RFs



(d) Individual estimates and CIs in RFs

Figure S.7: Evaluation of the nonparametric logistic model estimator. Figure S.7a and S.7c show the ROC curve with both AUC of 0.84. Figure S.7b and S.7d display estimated subject-level probabilities of ICU readmission and confidence intervals, illustrating heteroskedasticity across individuals.

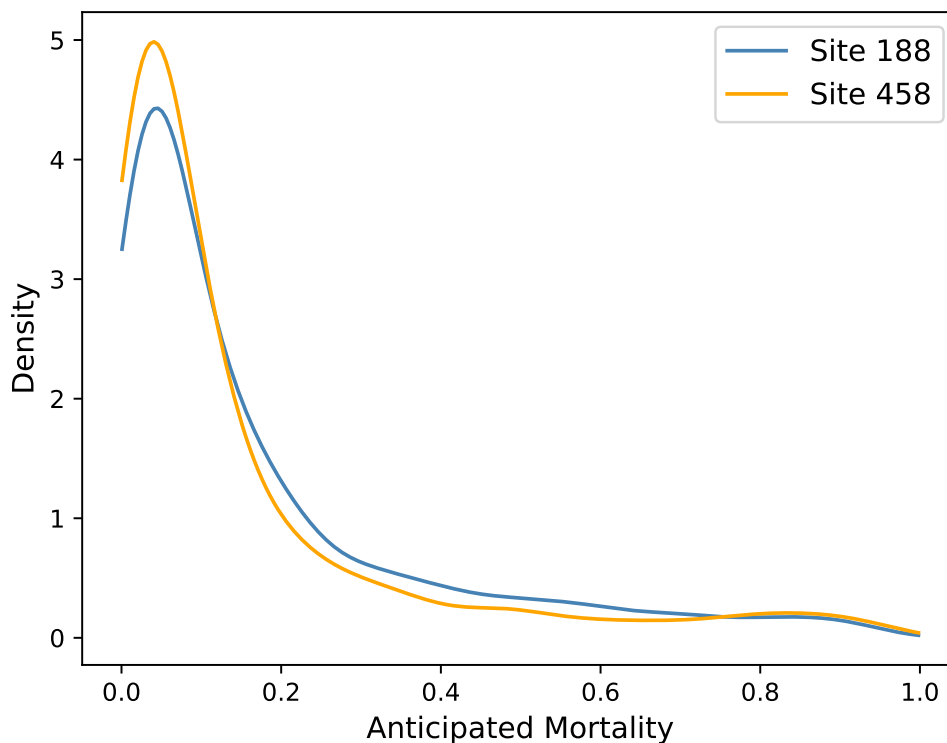


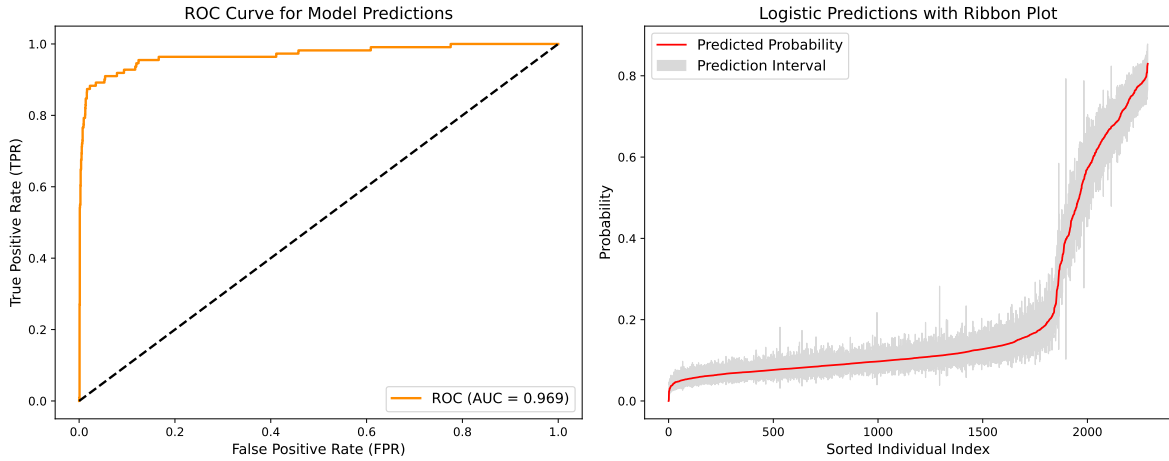
Figure S.8: Density of the anticipated mortality.

scenario for model generalization.

As shown in Figure S.9, the transferred model maintains strong predictive performance, achieving a high AUC. In the logistic regression, the true positive rate is 0.92 and the false positive rate is 0.10 at a 0.5 threshold. Notably, the model continues to produce heteroskedastic prediction intervals that vary across individuals, indicating that the ESM-based uncertainty quantification remains responsive to patient-level heterogeneity even under distributional shift. Although this evaluation is limited to a single external site, the results provide encouraging evidence that the proposed method generalizes effectively and that the predictive intervals meaningfully capture site-specific differences in risk. Broader multi-site validation is warranted.

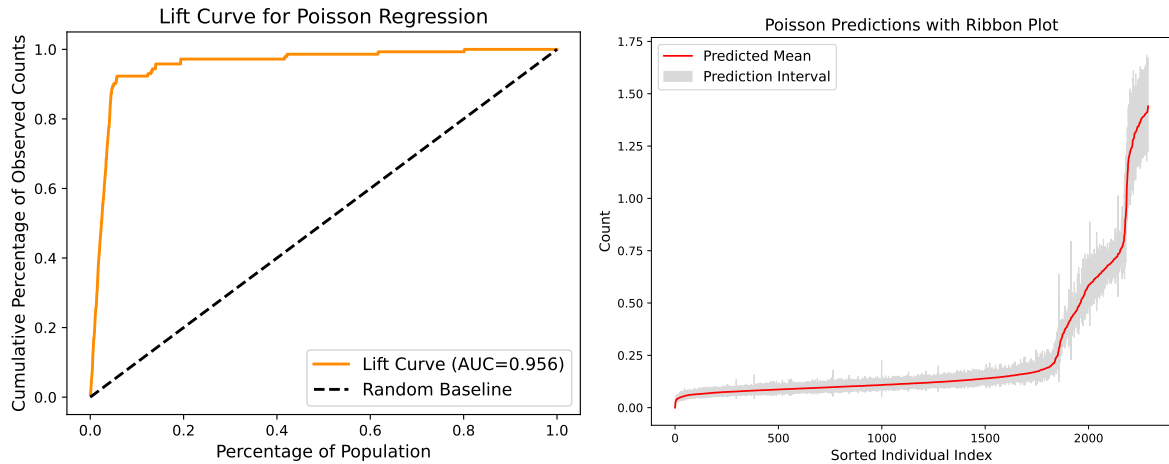
We also observed heteroskedasticity across individuals in the new site as well. Specifically, the variation in interval widths primarily stems from the input-dependent predictive

variance intrinsic to the ESM framework. In ESM, each prediction $\hat{f}^B(\mathbf{x}_*)$ aggregates outputs from multiple subsampled models, leading to ensemble variance that naturally differs across regions of the feature space. Consequently, the estimated uncertainty adapts to local variability, yielding a heteroskedastic characterization of estimation uncertainty.



(a) ROC curve with Logistic

(b) Individual estimates and CIs with Logistic



(c) Lift curve with Poisson

(d) Individual estimates and CIs with Poisson

Figure S.9: Evaluation of the transferred model from site 188 to site 458. Figure S.9a and S.9b show the results under logistic regression. Figure S.9c and S.9d display results with poisson regression.

Table S.17: Checklist of Some Notations Used in the Paper.

Symbol	Description	Symbol	Description
\mathbf{x}_i	Input covariate for the i -th observation	$\mathbb{E}[\cdot]$	Expectation operator
y_i	Response variable	$\text{Var}[\cdot]$	Variance operator
\mathbf{z}_i	(\mathbf{x}_i, y_i)	V_k	$\text{Var}(T_k(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}))$
B	Number of subsampling times	$J_{b_j i}$	$I(i \in \mathcal{I}^{b_j})$ indication function of i in subset b_j
$\hat{f}^b(\mathbf{x})$	Estimator for each subsample	$J_{\cdot i}$	$\sum_{j=1}^B J_{b_j i} / B$, average inclusion rate of sample i across B subsamples
$\hat{f}^B(\mathbf{x})$	Ensemble estimator	$Z_{b_j i}$	$(J_{b_j i} - J_{\cdot i})(\hat{f}^{b_j}(\mathbf{x}_*) - \hat{f}^B(\mathbf{x}_*))$
$\hat{f}^B(\mathbf{x}_*)$	$\sum_{i=1}^n (\mathbb{E}(\hat{f}^B(\mathbf{x}_*) \mathbf{z}_i) - \mathbb{E}\hat{f}^b(\mathbf{x}_*))$	M_i	$\mathbb{E}(Z_{b_j i} \mathbf{z}_1, \dots, \mathbf{z}_n)$, expectation takes over randomness of subsampling
$\xi_{1,r}(\mathbf{x}_*)$	$\text{Cov}(\hat{f}(\mathbf{x}_*; \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r), \hat{f}(\mathbf{x}_*, \mathbf{z}'_1, \mathbf{z}'_2, \dots, \mathbf{z}'_r))$	$\varepsilon_{b_j i}$	$Z_{b_j i} - M_i$, $\bar{\varepsilon}_i = \sum_{j=1}^B \varepsilon_{b_j i} / B$
$g_{1,r}(\mathbf{z}; \mathbf{x}_*)$	$\mathbb{E}\hat{f}^b(\mathbf{x}_*; \mathbf{z}, \mathbf{z}_2, \dots, \mathbf{z}_r) - \mathbb{E}\hat{f}^b(\mathbf{x}_*)$	\hat{V}_i	$\sum_{j=1}^B Z_{b_j i} / B$
$T_k(\mathbf{z}_1, \dots, \mathbf{z}_k)$	Hoeffding decomposition in (C.5)	$\widehat{\text{Var}}_i$	$\sum_{j=1}^B (Z_{b_j i} - \hat{V}_i)^2 / (B - 1)$
$W_{i_1 i_2 \dots i_k}$	Number of subsamples including i_1, \dots, i_k	A_i	$(\frac{r}{n} - \frac{r^2}{n^2})T_1(\mathbf{z}_i) + (\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2}) \sum_{j \neq i}^n T_1(\mathbf{z}_j)$

C Additional Propositions and Lemmas and Proofs of Theorems

For ease of reference, Table S.17 provides the key notation and glossary used throughout the paper, which will facilitate the technical developments that follow.

C.1 Preliminaries

We present the following proposition and lemmas, which are used to prove Theorem 3.5.

Let $\mathcal{N}_n(\tilde{\delta}, \mathcal{F}, \|\cdot\|_\infty)$ denote the covering number, i.e., the minimal number of $\|\cdot\|_\infty$ -balls with radius $\tilde{\delta}$ required to cover $\mathcal{F}(L, \mathbf{p}, s, F)$. For brevity, we write $\mathcal{N}_n = \mathcal{N}_n(\tilde{\delta}, \mathcal{F}, \|\cdot\|_\infty)$.

Then, for any $\hat{f}_n \in \mathcal{F}(L, \mathbf{p}, s, F)$, we have the following result.

Proposition C.1. *Under the conditions of Theorem 3.5, there exist constants $c', C' > 0$ depending only on κ , the function ψ , F , and K , such that*

$$\begin{aligned} \frac{1}{2}\Delta_n(\hat{f}_n) - c' \left(\frac{\log \mathcal{N}_n}{n} + \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + 1 \right) \tilde{\delta} \right) &\leq R_n(\hat{f}_n, f_0) \leq 2\Delta_n(\hat{f}_n) \\ + 2 \inf_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \mathbb{E}\ell(\mathbf{X}; f, f_0) + C' \left(\frac{\log \mathcal{N}_n}{n} + \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + 1 \right) \tilde{\delta} \right). \end{aligned}$$

The proof is provided in Section D. To obtain the bounds in Proposition C.1, we apply truncation techniques that account for heterogeneity across individuals, paving the way for establishing convergence results under generalized nonparametric regression models. The following two lemmas provide properties of the covering number and characterize the minimal distance between functions $f \in \mathcal{F}(L, \mathbf{p}, s, F)$ and f_0 .

Lemma C.2. *Let \mathcal{N}_n denote the covering number, i.e., the minimal number of $\|\cdot\|_\infty$ -balls*

with radius $\tilde{\delta}$ required to cover $\mathcal{F}(L, \mathbf{p}, s, \infty)$. Then

$$\log(\mathcal{N}_n) \leq (s+1) \log \left(2\tilde{\delta}^{-1}(L+1) \left(\prod_{l=0}^{L+1} (1+p_l) \right)^2 \right).$$

Lemma C.3. *Under the conditions of Theorem 3.5, it holds that*

$$\inf_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \|f - f_0\|_\infty^2 \lesssim \phi_n.$$

Proofs of Lemmas C.2 and C.3. See Lemma 5 and the proof of Theorem 1 in [Schmidt-Hieber \(2020\)](#) for the proofs of Lemmas C.2 and C.3. \square

The following lemma establishes key properties of the function ψ in the context of generalized nonparametric regression models.

Lemma C.4. *Under Assumption 3.1, there exist constants $c', C', C_{\text{Lip}} > 0$ such that for any $x, m \in [-C, C]$, where $C > 0$,*

$$\frac{C'}{2}(x-m)^2 \geq \psi(x) - \psi(m) - \psi'(m)(x-m) \geq \frac{c'}{2}(x-m)^2, \quad \psi(x) - \psi(m) \leq C_{\text{Lip}}|x-m|.$$

C.2 Proof of Theorem 3.5

With Lemma C.2 and if we let $\tilde{\delta} = n^{-1}$, it holds that

$$\log(\mathcal{N}_n) \leq (s+1) \log \left(2n(L+1)(\max\{p_l\} + 1)^{2L+4} \right).$$

Combining this inequality with Assumption 3.2, we have $s \lesssim n\phi_n \log n$, $\max\{p_l\} \lesssim n$, and the upper bound of $\log(\mathcal{N}_n)$ can be written as

$$\log(\mathcal{N}_n) \lesssim n\phi_n L \log^2 n. \tag{C.1}$$

By Proposition C.1, the lower bound for $R_n(\hat{f}_n, f_0)$ can be written as

$$\begin{aligned} \frac{1}{2}\Delta_n(\hat{f}_n) - R_n(\hat{f}_n, f_0) &\leq c' \left(\frac{\log \mathcal{N}_n}{n} + \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + 1 \right) n^{-1} \right) \\ &\lesssim \phi_n L \log^2 n. \end{aligned}$$

On the other hand, for the upper bound of $R_n(\widehat{f}_n, f_0)$, Lemma C.3 yields

$$\inf_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \|f - f_0\|_\infty^2 \lesssim \phi_n.$$

From Lemma C.4, it follows that $\mathbb{E} \ell(\mathbf{X}; f, f_0) \lesssim \|f - f_0\|_\infty^2$. Hence, combining the upper bound of $\log(\mathcal{N}_n)$ in (C.1) and $\inf_{f \in \mathcal{F}(L, \mathbf{p}, s, F)} \|f - f_0\|_\infty^2$ above, we have

$$R_n(\widehat{f}_n, f_0) - 2\Delta_n(\widehat{f}_n) \lesssim \phi_n L \log^2 n.$$

C.3 Proof of Proposition 3.6

We adapt the proof technique in Schmidt-Hieber (2020) to construct the approximation error in such class of ReLU networks. Suppose that

$$\phi_r = r^{-\frac{2m}{2m+t}},$$

and let $c = \frac{2m}{2m+t}$. Then we specifically construct the neural network \tilde{f} with the approximation rate

$$\sup_{f_0 \in \mathcal{G}} \|\tilde{f} - f_0\|_\infty \lesssim \phi_r^{\frac{1+\delta(t+2m)/t}{2}}.$$

As we change the sparsity condition slightly, this rate can be proved to be much faster than the prediction error $\phi_r^{(1+\delta)/2}$ with the assumption $s \asymp r^{1+\delta} \phi_r \log r$, which completes the proof of Proposition 3.6.

C.4 Proof of Theorem 3.7

By Assumption 3.4, recall $\gamma > \frac{1}{1 + \min_{j=0, \dots, q} (2m_j^*) / (2m_j^* + t_j)}$. For the sufficiently small $\varepsilon > 0$, it is easy to see that

$$\lim_{n \rightarrow \infty} n \phi_r L \log^2(r) / (r^2 \inf_{\mathbf{x} \in \mathcal{X}} \xi_{1,r}(\mathbf{x})) = 0,$$

there always exists a positive sequence δ_n with $\delta_n \rightarrow 0$ such that

$$\lim_{n \rightarrow \infty} n\phi_r L \log^2(r) / (\delta_n r^2 \inf_{\mathbf{x} \in \mathcal{X}} \xi_{1,r}(\mathbf{x})) = 0.$$

For such δ_n , we construct the set \mathcal{A}_{δ_n} as

$$\mathcal{A}_{\delta_n} = \left\{ \mathbf{X} \in \mathcal{X} : [\mathbb{E}\widehat{f}(\mathbf{X}) - f_0(\mathbf{X})|\mathbf{X}]^2 \leq \frac{C\phi_r L \log^2 r}{\delta_n} \right\},$$

By Chebyshev's inequality we have

$$\begin{aligned} \mathbb{P}_{\mathbf{x}}(\mathcal{A}_{\delta_n}) &= 1 - \mathbb{P}_{\mathbf{x}}\left([\mathbb{E}\widehat{f}(\mathbf{X}) - f_0(\mathbf{X})|\mathbf{X}]^2 \geq \frac{C\phi_r L \log^2 r}{\delta_n}\right) \\ &\geq 1 - \delta_n. \end{aligned}$$

Moreover, by the definition of \mathcal{A}_{δ_n} , we have for every fixed $\mathbf{x}_* \in \mathcal{A}_{\delta_n}$,

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} |\mathbb{E}(\widehat{f}^{b_j}(\mathbf{x}_*) - f_0(\mathbf{x}_*))| = O\left(\sqrt{\frac{n\phi_r L \log^2 r}{\delta_n r^2 \xi_{1,r}(\mathbf{x}_*)}}\right) \rightarrow 0,$$

where the last equality is by the condition $\lim_{n \rightarrow \infty} n\phi_r L \log^2(r) / (\delta_n r^2 \inf_{\mathbf{x} \in \mathcal{X}} \xi_{1,r}(\mathbf{x})) = 0$.

By the basic inequality of $|\mathbb{E}\widehat{f}^B(\mathbf{x}_*) - f_0(\mathbf{x}_*)|^2 \leq \sum_{j=1}^B (\mathbb{E}\widehat{f}^{b_j}(\mathbf{x}_*) - f_0(\mathbf{x}_*))^2 / B$, we have that

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} |\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - f_0(\mathbf{x}_*))| = O\left(\sqrt{\frac{n\phi_r L \log^2 r}{\delta_n r^2 \xi_{1,r}(\mathbf{x}_*)}}\right) \rightarrow 0 \quad (\text{C.2})$$

for any $\mathbf{x}_* \in \mathcal{A}_{\delta_n}$.

Assume the observed data of size n , $\mathcal{D}_n = \{(y_i, \mathbf{x}_i) : i = 1, \dots, n\}$, are independently and identically distributed copies of (y, \mathbf{X}) , with $y \in \mathbb{R}^1$ and $\mathbf{X} \in \mathbb{R}^{d \times 1}$. Let $\mathbf{z}_i = (y_i, \mathbf{x}_i)$ represent an independent observation of sample points, where y_i is the label and \mathbf{x}_i is the response. By the definition of $\xi_{1,r}(\mathbf{x})$, it is clear that

$$\xi_{1,r}(\mathbf{x}_*) = \text{Cov}(\widehat{f}(\mathbf{x}_*; \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r), \widehat{f}(\mathbf{x}_*; \mathbf{z}'_1, \mathbf{z}'_2, \dots, \mathbf{z}'_r)).$$

Here, $\widehat{f}(\mathbf{x}_*; \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r)$ means that we obtain \widehat{f} from the subsample $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_r$, and then apply it to the point \mathbf{x}_* . The terms \mathbf{z}'_i and \mathbf{z}_i are independently generated from the same data generation process. In the following of the proofs, all the expectations and variance are taken over the fixed \mathbf{x}_* .

Next, we consider the asymptotic normality of the first order dominants of the Hoeffding decomposition. By the definition of $\widehat{f}^B(\mathbf{x}_*)$, it is clear that \widehat{f}^{b_j} and $\widehat{f}^{b_{j'}}$ are not independent given \mathcal{I}^{b_j} and $\mathcal{I}^{b_{j'}}$ have overlap. To handle the dependency, we consider the Hoeffding decomposition. For the n sample data points $\mathbf{z}_i = (y_i, \mathbf{x}_i)$, define

$$\mathring{f}^B(\mathbf{x}_*) = \sum_{i=1}^n \left(\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) | \mathbf{z}_i) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) \right).$$

Here, the expectation \mathbb{E} is taken over the training sample. By the definition of $\widehat{f}^B(\mathbf{x}_*)$, we have

$$\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - (\mathbb{E}\widehat{f}^b(\mathbf{x}_*)) | \mathbf{z}_i) = \mathbb{E} \left(\frac{1}{B} \sum_{j=1}^B \widehat{f}^{b_j}(\mathbf{x}_*) | \mathbf{z}_i \right) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*).$$

Let W_{b_j} be the number of subsamples \mathcal{I}^{b_j} that contain index i , and define $g_{1,r}(\mathbf{z}; \mathbf{x}_*) = \mathbb{E}\widehat{f}^b(\mathbf{x}_*; \mathbf{z}, \mathbf{z}_2, \dots, \mathbf{z}_r) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*)$. We have W_{b_j} is independent of \mathbf{z}_i and $\mathbb{E}W_{b_j} = B \binom{n-1}{r-1} / \binom{n}{r}$.

Hence, the equation above involving the term \mathbf{z}_i can be further expressed as

$$\begin{aligned} & \mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - (\mathbb{E}\widehat{f}^b(\mathbf{x}_*)) | \mathbf{z}_i) \\ &= \frac{1}{B} \sum_{j=1}^B \mathbb{E}_{W_{b_j}} \left(\mathbb{E}(\widehat{f}^{b_j}(\mathbf{x}_*) - (\mathbb{E}\widehat{f}^b(\mathbf{x}_*)) | \mathbf{z}_i, W_{b_j}) \right) \\ &= \frac{1}{B} \mathbb{E}_{W_{b_j}} W_{b_j} g_{1,r}(\mathbf{z}_i; \mathbf{x}_*) + \frac{1}{B} \mathbb{E}_{W_{b_j}} (B - W_{b_j}) \cdot \mathbb{E}(\widehat{f}^{b_j}(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*)) \\ &= \frac{r}{n} g_{1,r}(\mathbf{z}_i; \mathbf{x}_*). \end{aligned}$$

Therefore we conclude that

$$\mathring{f}^B(\mathbf{x}_*) = \frac{r}{n} \sum_{i=1}^n g_{1,r}(\mathbf{z}_i; \mathbf{x}_*). \quad (\text{C.3})$$

Equation (C.3) provides a nice form, as $g_{1,r}(\mathbf{z}_i)$ is independent of each other. Moreover, by the assumption that \hat{f}^b is bounded, the Lindeberg condition is satisfied, and it is clear that

$$\frac{\sqrt{n}}{r \cdot \sqrt{\text{Var}(g_{1,r}(\mathbf{z}_i; \mathbf{x}_*))}} \cdot \mathring{f}^B(\mathbf{x}_*) \xrightarrow{d} \mathcal{N}(0, 1).$$

By the definition of $\xi_{1,r}(\mathbf{x}_*)$, it is easy to see that $\text{Var}(g_{1,r}(\mathbf{z}_i; \mathbf{x}_*)) = \xi_{1,r}(\mathbf{x}_*)$, hence we have

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot \mathring{f}^B(\mathbf{x}_*) \xrightarrow{d} \mathcal{N}(0, 1). \quad (\text{C.4})$$

In the next step, we would like to substitute $\mathring{f}^B(\mathbf{x}_*)$ with $\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*)$. Applying the Hoeffding decomposition, it is clear that the simple $\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r)$ can be written as

$$\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) = \sum_{i=1}^r T_1(\mathbf{z}_i) + \sum_{i<j} T_2(\mathbf{z}_i, \mathbf{z}_j) \dots + T_r(\mathbf{z}_1, \dots, \mathbf{z}_r). \quad (\text{C.5})$$

The $2^r - 1$ random variables on the right hand side are all mean-zero and uncorrelated. Here, $T_1(\mathbf{z}_i) = \mathbb{E}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r) | \mathbf{z}_i) - \mathbb{E}\hat{f}^b(\mathbf{x}_*)$. Applying the decomposition to $\hat{f}^B(\mathbf{x}_*)$ and utilizing the same trick we obtain (C.3), we have

$$\begin{aligned} \mathbb{E}(\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) | \mathbf{z}_1, \dots, \mathbf{z}_n) &= \frac{1}{\binom{n}{r}} \left(\binom{n-1}{r-1} \sum_{i=1}^n T_1(\mathbf{z}_i) + \binom{n-2}{r-2} \sum_{i<j} T_2(\mathbf{z}_i, \mathbf{z}_j) \right. \\ &\quad \left. + \dots + \sum_{i_1 < \dots < i_r} T_r(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_r}) \right). \end{aligned}$$

Hence, it is clear that $\mathring{f}^B(\mathbf{x}_*)$ is exactly the first order in the Hoeffding decomposition. We have

$$\begin{aligned} &\mathbb{E}(\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*) | \mathbf{z}_1, \dots, \mathbf{z}_n) \\ &= \frac{1}{\binom{n}{r}} \left(\binom{n-2}{r-2} \sum_{i<j} T_2(\mathbf{z}_i, \mathbf{z}_j) + \dots + \sum_{i_1 < \dots < i_r} T_r(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_r}) \right). \end{aligned}$$

Denote by $V_k = \text{Var}(T_k(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}))$, we then calculate the value of $\mathbb{E}(\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2$, which we will show later this term is negligible. To calculate this term, we apply

the equality $\mathbb{E}X_n^2 = \mathbb{E}\text{Var}(X_n|Y_n) + \mathbb{E}(\mathbb{E}X_n|Y_n)^2$. Then the term $\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2$ can be written as

$$\begin{aligned} \mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2 &= \underbrace{\mathbb{E}(\text{Var}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)|\mathbf{z}_1, \dots, \mathbf{z}_n))}_{J_1} \\ &\quad + \underbrace{\mathbb{E}(\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)|\mathbf{z}_1, \dots, \mathbf{z}_n))^2}_{J_2}. \end{aligned}$$

Here, we let $X_n = \widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)$ and $Y_n = \mathbf{z}_1, \dots, \mathbf{z}_n$. From the expression of $\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)|\mathbf{z}_1, \dots, \mathbf{z}_n)$ above, we can easily see that

$$\begin{aligned} J_2 &= \frac{1}{\binom{n}{r}^2} \cdot \sum_{k=2}^r \binom{n-k}{r-k}^2 \cdot \binom{n}{k} V_k \\ &= \sum_{k=2}^r \frac{(n-k)!r!}{n!(r-k)!} \binom{r}{k} V_k \\ &\leq \frac{r(r-1)}{n(n-1)} \cdot \sum_{k=2}^r \binom{r}{k} V_k \\ &\leq \frac{r(r-1)}{n(n-1)} \text{Var}(\widehat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r)). \end{aligned}$$

Here, the first equality is by the fact that all the terms are uncorrelated, the first inequality is by the decreasing monotonicity of $\frac{(n-k)!r!}{n!(r-k)!}$ with respect to k , and the last inequality is by the decomposition of \widehat{f}^b in (C.5). To calculate $\mathbb{E}(\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2$, it remains for us to calculate J_1 . Let $W_{i_1 i_2 \dots i_k}$ ($k \leq r$) be the number of the B subsamples that contain index i_1, \dots, i_k simultaneously. Then we have

$$\widehat{f}^B(\mathbf{x}_*) - \mathbb{E}\widehat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*) = \frac{1}{B} \left(\sum_{i_1 < i_2}^n W_{i_1 i_2} T_2(\mathbf{z}_{i_1}, \mathbf{z}_{i_2}) + \dots + \sum_{i_1 < \dots < i_r}^n W_{i_1 \dots i_r} T_r(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_r}) \right).$$

By the independence on the different terms T_k , we conclude that

$$\begin{aligned}
J_1 &= \frac{1}{B^2} \sum_{k=2}^r \text{Var}(W_{i_1 \dots i_k}) \binom{n}{k} V_k \\
&\leq \frac{1}{B} \sum_{k=2}^r \frac{\binom{n-k}{r-k}}{\binom{n}{r}} \binom{n}{k} V_k \\
&= \frac{1}{B} \sum_{k=2}^r \binom{r}{k} V_k \\
&\leq \frac{\text{Var}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r))}{B}.
\end{aligned}$$

Here, the first inequality is by $\text{Var}(W_{i_1 \dots i_k}) \leq B \cdot \frac{\binom{n-k}{r-k}}{\binom{n}{r}}$ and the uncorrelation on different terms T_k , and the last inequality is by $\sum_{k=2}^r \binom{r}{k} V_k \leq \text{Var}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r))$. We hence conclude that

$$\begin{aligned}
\mathbb{E}(\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2 &= J_1 + J_2 \\
&\leq \left(\frac{r(r-1)}{n(n-1)} + \frac{1}{B} \right) \text{Var}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r)),
\end{aligned}$$

leading to

$$\begin{aligned}
\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \mathbb{E}(\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*))^2 &\leq \left(\frac{1}{n \xi_{1,r}(\mathbf{x}_*)} + \frac{n}{B r^2 \xi_{1,r}(\mathbf{x}_*)} \right) \text{Var}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r)) \\
&\xrightarrow{p} 0.
\end{aligned} \tag{C.6}$$

Here, we utilize the fact that $\hat{f}^b(\mathbf{x}_*)$ is bounded. The bound of $\hat{f}^b(\mathbf{x}_*)$ indicates that $\text{Var}(\hat{f}^b(\mathbf{x}_*; \mathbf{z}_1, \dots, \mathbf{z}_r))$ is bounded. Additionally, we rely on Assumption 3.4 that $\liminf_{n \rightarrow +\infty} n \xi_{1,r}(\mathbf{x}_*) \geq n^{1-\varepsilon-\gamma} \rightarrow +\infty$ and $\liminf_{n \rightarrow +\infty} B r^2 \xi_{1,r}(\mathbf{x}_*)/n \rightarrow +\infty$. From (C.6), we have

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot (\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)) \xrightarrow{p} 0.$$

Therefore it holds that

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot (\hat{f}^B(\mathbf{x}_*) - \mathbb{E}\hat{f}^b(\mathbf{x}_*) - \mathring{f}^B(\mathbf{x}_*)) \xrightarrow{p} 0.$$

From (C.4), we also have that

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot \hat{f}^B(\mathbf{x}_*) \xrightarrow{p} \mathcal{N}(0, 1),$$

Combing the convergence above, we have

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot (\hat{f}^B(\mathbf{x}_*) - \mathbb{E} \hat{f}^b(\mathbf{x}_*)) \xrightarrow{d} \mathcal{N}(0, 1).$$

The only thing that needs to prove is

$$\sqrt{\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)}} \cdot (f_0(\mathbf{x}_*) - \mathbb{E} \hat{f}^b(\mathbf{x}_*)) \rightarrow 0$$

for $\mathbf{x}_* \in \mathcal{A}_{\delta_n}$. This is proved by (C.2) and $\mathbb{E} \hat{f}^B(\mathbf{x}_*) = \mathbb{E} \hat{f}^b(\mathbf{x}_*)$. Wrapping all together, we complete the proof of Theorem 3.7.

C.5 Proof of Theorem 3.8

Recall the definition

$$Z_{b_{ji}} = (J_{b_{ji}} - J_{\cdot i})(\hat{f}^{b_{ji}}(\mathbf{x}_*) - \hat{f}^B(\mathbf{x}_*)), \quad \hat{V}_i = \frac{\sum_{j=1}^B Z_{b_{ji}}}{B}.$$

We denote by

$$\tilde{\sigma}^2 = \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \hat{V}_i^2.$$

Without loss of generality, we calculate the term \hat{V}_1 , and the terms \hat{V}_i with $i \neq 1$ are exactly the same. Consider \hat{V}_1^2 and $B/n \rightarrow c$ for some constant $c > 0$. Then by the inspiration of the proof of Lemma 13 of Wager & Athey (2018), we can define

$$\begin{aligned} M_1 &= \left(\frac{r}{n} - \frac{r^2}{n^2} \right) T_1(\mathbf{z}_1) + \left(\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2} \right) \sum_{i=2}^n T_1(\mathbf{z}_i) \\ &+ \frac{r}{n} \sum_{k=1}^{r-1} \left(\frac{\binom{r-1}{k}}{\binom{n-1}{k}} - \frac{\binom{r}{k+1}}{\binom{n}{k+1}} \right) \sum_{2 \leq i_1 < \dots < i_k \leq n} T_{k+1}(\mathbf{z}_1, \mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}) \\ &+ \frac{r}{n} \sum_{k=2}^r \left(\frac{\binom{r-1}{k}}{\binom{n-1}{k}} - \frac{\binom{r}{k}}{\binom{n}{k}} \right) \sum_{2 \leq i_1 < \dots < i_k \leq n} T_k(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}), \end{aligned}$$

and M_i the same way with index 1 of \mathbf{z} be replaced by i of \mathbf{z} . Let the whole data set $\mathcal{D}_n = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$, and recall $Z_{b_j i} = (J_{b_j i} - J_{\cdot i})(\hat{f}^{b_j}(\mathbf{x}_*) - \hat{f}^B(\mathbf{x}_*))$, the motivation of the definition of M_i is given by the direct calculation that

$$\mathbb{E}(Z_{b_j i} | \mathcal{D}_n) = M_i.$$

Here, the expectation takes over the randomness of subsampling. Define

$$\hat{\sigma}_1^2 = \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n M_i^2, \quad R_B = \tilde{\sigma}^2 - \hat{\sigma}_1^2.$$

We next prove two key conclusions, which will show the consistency of our results.

$$\begin{aligned} \frac{n(n-1)}{(n-r)^2} \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B (Z_{b_j i} - \hat{V}_i)^2 - R_B &= o_p\left(\frac{r^2 \xi_{1,r}(\mathbf{x}_*)}{n}\right), \\ \sqrt{\frac{r^2 \xi_{1,r}(\mathbf{x}_*)}{n \hat{\sigma}_1^2}} &\xrightarrow{p} 1, \end{aligned} \tag{C.7}$$

In the following of the proof, we denote by $\hat{R}_B = \frac{n(n-1)}{(n-r)^2} \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B (Z_{b_j i} - \hat{V}_i)^2$. If

(C.7) holds, we can see that

$$\frac{n \hat{\sigma}_*^2}{r^2 \xi_{1,r}(\mathbf{x}_*)} = \frac{n(\tilde{\sigma}^2 - \hat{R}_B)}{r^2 \xi_{1,r}(\mathbf{x}_*)} = \frac{n(R_B - \hat{R}_B)}{r^2 \xi_{1,r}(\mathbf{x}_*)} + \frac{n(\tilde{\sigma}^2 - R_B)}{r^2 \xi_{1,r}(\mathbf{x}_*)} = o_p(1) + \frac{n \hat{\sigma}_1^2}{r^2 \xi_{1,r}(\mathbf{x}_*)} \xrightarrow{p} 1,$$

which completes the proof of Theorem 3.8. It remains for us to prove (C.7).

We first investigate the term $\hat{R}_B - R_B = \frac{n(n-1)}{(n-r)^2} \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B (Z_{b_j i} - \hat{V}_i)^2 - R_B$ in (C.7). Recall $R_B = \tilde{\sigma}^2 - \hat{\sigma}_1^2$, it is easy to see that

$$R_B = \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n (\hat{V}_i^2 - M_i^2).$$

Let $\varepsilon_{b_j i} = Z_{b_j i} - M_i$, and $\bar{\varepsilon}_i = \frac{1}{B} \sum_{j=1}^B \varepsilon_{b_j i}$, recall $\hat{V}_i = \frac{\sum_{j=1}^B Z_{b_j i}}{B}$, we have that $\hat{V}_i = M_i + \bar{\varepsilon}_i$ and $\hat{V}_i^2 - M_i^2 = 2M_i \bar{\varepsilon}_i + \bar{\varepsilon}_i^2$, which means

$$R_B = \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n 2M_i \bar{\varepsilon}_i + \bar{\varepsilon}_i^2.$$

Further denote by $\widehat{\text{Var}}_i = \frac{1}{B-1} \sum_{j=1}^B (Z_{b_j i} - \widehat{V}_i)^2$, we have

$$\widehat{R}_B = \frac{n(n-1)}{(n-r)^2} \frac{1}{B} \sum_{i=1}^n \widehat{\text{Var}}_i.$$

We then combine the expression of \widehat{R}_B and R_B above and get that

$$R_B - \widehat{R}_B = \frac{n(n-1)}{(n-r)^2} \left(\underbrace{\sum_{i=1}^n 2M_i \bar{\varepsilon}_i}_{\Delta_1} + \underbrace{\sum_{i=1}^n \bar{\varepsilon}_i^2 - B^{-1} \text{Var}_i}_{\Delta_2} - \underbrace{\frac{1}{B} \sum_{i=1}^n \widehat{\text{Var}}_i - \text{Var}_i}_{\Delta_3} \right).$$

Here, $\text{Var}_i = \text{Var}_b(Z_{b_j i})$ with variance condition on the dataset \mathcal{D}_n . We bound the three terms one by one. For the term Δ_1 , by the definition of M_i and $\bar{\varepsilon}_i$, we have

$$\Delta_1 = \frac{2}{B} \sum_{j=1}^B \sum_{i=1}^n M_i \varepsilon_{b_j i}.$$

Conditional on the dataset \mathcal{D}_n , $\sum_{i=1}^n M_i \varepsilon_{b_j i}$ is independent with different b_j . It is easy to see that $\mathbb{E}_b \Delta_1 = 0$, and we give the calculation of $\text{Var}_b \Delta_1$. From the expression of Δ_1 , conditional on data set \mathcal{D}_n , $\varepsilon_{b_j i}$ are independent for different b_j , it is easy to see that

$$\text{Var}_b \Delta_1 \asymp \frac{\text{Var}_b \sum_{i=1}^n M_i \varepsilon_{b_j i}}{B}.$$

It remains for us to investigate the term in the numerator to bound Δ_1 . We have

$$\text{Var}_b \sum_{i=1}^n M_i \varepsilon_{b_j i} = \sum_{i=1}^n M_i^2 \text{Var}_b \varepsilon_{b_j i} + \sum_{i \neq i'} M_i M_{i'} \text{Cov}_b(\varepsilon_{b_j i}, \varepsilon_{b_j i'}).$$

From the definition of M_i , we can easily bound $M_i = O(r/n)$, for the covariance of $\varepsilon_{b_j i}$ and $\varepsilon_{b_j i'}$, simple calculation of the subsampling randomness gives us that $\text{Var}_b \varepsilon_{b_j i} = O(\frac{r}{n}(1 - \frac{r}{n})) = O(r/n)$, and $\text{Cov}_b(\varepsilon_{b_j i}, \varepsilon_{b_j i'}) = O(|\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2}|) = O(\frac{r}{n^2})$, we hence have

$$\begin{aligned} \text{Var}_b \sum_{i=1}^n M_i \varepsilon_{b_j i} &= O\left(n \cdot \frac{r^2}{n^2} \cdot \frac{r}{n} + n^2 \cdot \frac{r^2}{n^2} \cdot \frac{r}{n^2}\right) \\ &= O\left(\frac{r^3}{n^2}\right). \end{aligned}$$

This indicates that $\text{Var}_b \Delta_1 = O(\frac{r^3}{Bn^2})$. Combining the fact that $\mathbb{E}_b \Delta_1 = 0$, we conclude that $|\Delta_1| = O_p(\frac{r^{3/2}}{n\sqrt{B}})$.

For the term Δ_2 , by the expression that $\bar{\varepsilon}_i^2 = \frac{1}{B^2} \sum_{j=1}^B \varepsilon_{b_j i}^2 + \frac{1}{B^2} \sum_{j \neq j'} \varepsilon_{b_j i} \varepsilon_{b_{j'} i}$, Δ_2 can be further written as

$$\Delta_2 = \underbrace{\frac{1}{B^2} \sum_{i=1}^n \sum_{j=1}^B (\varepsilon_{b_j i}^2)}_{\Delta_{21}} - B^{-1} \text{Var}_i + \underbrace{\frac{1}{B^2} \sum_{i=1}^n \sum_{j \neq j'} \varepsilon_{b_j i} \varepsilon_{b_{j'} i}}_{\Delta_{22}}.$$

We have $\mathbb{E}_b \Delta_{21} = \mathbb{E}_b \Delta_{22} = 0$. Simple calculation based on the randomness of subsampling gives $\text{Var}_b(\Delta_{21}) = O(n^2/B^3)$, we have $|\Delta_{21}| = O_p(n/B^{3/2})$. As for Δ_{22} , we consider $T_{b_j, b_{j'}} = \sum_{i=1}^n \varepsilon_{b_j i} \varepsilon_{b_{j'} i}$, and we have $\mathbb{E}_b T_{b_j, b_{j'}}^2 = O(r^2/n)$, and $\mathbb{E}_b T_{b_j, b_{j'}} \cdot T_{b_{j'}, b_j} = 0$ as long as $b_j, b_{j'}, b'_j, b'_j$ has three or more distinct values by the conditional independence. We have

$$\text{Var}_b(\Delta_{22}) \leq \frac{1}{B^4} \cdot \mathbb{E}_b \sum_{b_j \neq b_{j'}, b'_j \neq b'_j} T_{b_j, b_{j'}} \cdot T_{b'_j, b'_j} = O\left(\frac{r^2}{nB^2}\right)$$

We conclude that $|\Delta_2| = O_p(n/B^{3/2} + r/(B\sqrt{n}))$.

For the term Δ_3 , by the definition of $\widehat{\text{Var}}_i = \frac{1}{B-1} \sum_{j=1}^B (Z_{b_j i} - \widehat{V}_i)^2$, it holds that

$$\begin{aligned} \Delta_3 &= \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B \left((Z_{b_j i} - \widehat{V}_i)^2 - \frac{B-1}{B} \text{Var}_i \right) \\ &= \left(\frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B \left((Z_{b_j i} - M_i)^2 - \text{Var}_i \right) \right) + O\left(\frac{n}{B^2}\right). \end{aligned}$$

Here, the first equality is by direct calculation and the second equality is by Var_i is bounded.

Let , easy to bound that $\text{Var}_b(T_{b_j}) = O(n^2)$, we have that

$$\begin{aligned} \mathbb{E}_b \left\{ \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B \left((Z_{b_j i} - M_i)^2 - \text{Var}_i \right) \right\} &= 0; \\ \text{Var}_b \left\{ \frac{1}{B(B-1)} \sum_{i=1}^n \sum_{j=1}^B \left((Z_{b_j i} - M_i)^2 - \text{Var}_i \right) \right\} &= O(n^2/B^3), \end{aligned}$$

we hence conclude that $|\Delta_3| = O_p(n/B^{3/2})$.

We collect all the error terms of Δ_1 , Δ_2 and Δ_3 , and have

$$\Delta_1 = O_p\left(\frac{r^{3/2}}{n\sqrt{B}}\right), \quad \Delta_2 = O_p(n/B^{3/2} + r/(B\sqrt{n})), \quad \Delta_3 = O_p(n/B^{3/2}),$$

easy from order comparison to see that if $\Delta_1 = o_p\left(\frac{r^2\xi_{1,r}(\mathbf{x}_*)}{n}\right)$, then $R_B - \widehat{R}_B = o_p\left(\frac{r^2\xi_{1,r}(\mathbf{x}_*)}{n}\right)$.

We have

$$\frac{n\Delta_1}{r^2\xi_{1,r}(\mathbf{x}_*)} = \frac{1}{\sqrt{rB}\xi_{1,r}(\mathbf{x}_*)} = o_p(1),$$

where we use the Assumption 3.4 that $B \gtrsim n$ and ε small enough. We hence conclude that $R_B - \widehat{R}_B = o_p\left(\frac{r^2\xi_{1,r}(\mathbf{x}_*)}{n}\right)$.

It remains for us to prove $\sqrt{\frac{r^2\xi_{1,r}(\mathbf{x}_*)}{n\widehat{\sigma}_1^2}} \xrightarrow{p} 1$. All the terms T_k on the right hand side in M_i are uncorrelated. We define $A_1 = \left(\frac{r}{n} - \frac{r^2}{n^2}\right)T_1(\mathbf{z}_1) + \left(\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2}\right)\sum_{i=2}^n T_1(\mathbf{z}_i)$ which is the first line of the equation above. We then have $\mathbb{E}(M_1 - A_1) \cdot A_1 = 0$. Similarly we define A_i with the replace of 1 by i . we give the calculation of three terms: $\mathbb{E}A_i \cdot T_1(\mathbf{z}_i)$, $\mathbb{E}A_i^2$ and $\mathbb{E}(M_i - A_i)^2$, and then give the proof based on these terms.

For the term $\mathbb{E}A_i \cdot T_1(\mathbf{z}_i)$, we have

$$\mathbb{E}A_i \cdot T_1(\mathbf{z}_i) = \mathbb{E}\left(\frac{r}{n} - \frac{r^2}{n^2}\right)T_1^2(\mathbf{z}_1) = \left(\frac{r}{n} - \frac{r^2}{n^2}\right)\xi_{1,r}(\mathbf{x}_*). \quad (\text{C.8})$$

For the term $\mathbb{E}A_i^2$, we have

$$\begin{aligned} \mathbb{E}A_i^2 &= \frac{r^2(n-r)^2}{n^4}\xi_{1,r}(\mathbf{x}_*) + \frac{r^2(n-1)}{n^2}\left(\frac{r-1}{n-1} - \frac{r}{n}\right)^2\xi_{1,r}(\mathbf{x}_*) \\ &= \frac{r^2(n-r)^2}{n^3(n-1)}\xi_{1,r}(\mathbf{x}_*). \end{aligned} \quad (\text{C.9})$$

For the term $\mathbb{E}(\widehat{V}_i - A_i)^2$, we have

$$\begin{aligned} \mathbb{E}(M_i - A_i)^2 &= \frac{r^2}{n^2}\sum_{k=1}^{r-1}\binom{n-1}{k}\left(\frac{\binom{r-1}{k}}{\binom{n-1}{k}} - \frac{\binom{r}{k+1}}{\binom{n}{k+1}}\right)^2V_{k+1} + \frac{r^2}{n^2}\sum_{k=2}^r\binom{n-1}{k}\left(\frac{\binom{r-1}{k}}{\binom{n-1}{k}} - \frac{\binom{r}{k}}{\binom{n}{k}}\right)^2V_k \\ &\lesssim \frac{r^2(n-1)}{n^2}\left(\frac{r-1}{n-1} - \frac{\binom{r}{2}}{\binom{n}{2}}\right)^2\binom{r}{2}^{-1}\text{Var}[\widehat{f}^b(\mathbf{x}_*)] \\ &\lesssim \frac{2r^2\text{Var}[\widehat{f}^b(\mathbf{x}_*)]}{n^3}. \end{aligned} \quad (\text{C.10})$$

Here, the first inequality is by the fact that summation in the first line is maximized in the second-order terms and $\sum_{k=1}^r\binom{r}{k}V_k = \text{Var}(\widehat{f}^b(\mathbf{x}_*)) \leq C < +\infty$, and the second inequality

is by simple calculation. We hence have

$$\begin{aligned}
& \frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \mathbb{E} \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \left(A_i - \frac{r(n-r)}{n^2} T_1(\mathbf{z}_i) \right)^2 \\
&= \frac{n^3(n-1)}{r^2(n-r)^2 \xi_{1,r}(\mathbf{x}_*)} \mathbb{E} \left(A_i - \frac{r(n-r)}{n^2} T_1(\mathbf{z}_i) \right)^2 \\
&= \frac{n^3(n-1)}{r^2(n-r)^2} \left(\frac{r^2(n-r)^2}{n^3(n-1)} - \frac{2r^2(n-r)^2}{n^4} + \frac{r^2(n-r)^2}{n^4} \right) = \frac{1}{n} \rightarrow 0.
\end{aligned}$$

Here, the last second equality comes from (C.8) and (C.9). For any sequence of random variable $X_n > 0$, $\mathbb{E}X_n \rightarrow 0$ implies $X_n \xrightarrow{p} 0$ by the Markov inequality. Therefore, we conclude that

$$\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \left(A_i - \frac{r(n-r)}{n^2} T_1(\mathbf{z}_i) \right)^2 \xrightarrow{p} 0. \quad (\text{C.11})$$

By the weak law of large number, we also have

$$\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \frac{r^2(n-r)^2}{n^4} T_1^2(\mathbf{z}_i) \xrightarrow{p} 1. \quad (\text{C.12})$$

Here $T_1^2(\mathbf{z}_i)$ are iid and the expectation $\mathbb{E}T_1^2(\mathbf{z}_i) = \xi_{1,r}(\mathbf{x}_*)$. We then investigate the convergence of the term

$$\frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \frac{r}{n} A_i \cdot T_1(\mathbf{z}_i).$$

Recalling the definition of $A_i = \left(\frac{r}{n} - \frac{r^2}{n^2}\right) T_1(\mathbf{z}_i) + \left(\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2}\right) \sum_{i' \neq i} T_1(\mathbf{z}_{i'})$, the equation above can be simplified and written as

$$\begin{aligned}
& \frac{n}{r^2 \xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \frac{r(n-r)}{n^2} A_i \cdot T_1(\mathbf{z}_i) \sim \frac{n}{r(n-r) \xi_{1,r}(\mathbf{x}_*)} \sum_{i=1}^n A_i \cdot T_1(\mathbf{z}_i) \\
&= \frac{n}{r(n-r) \xi_{1,r}(\mathbf{x}_*)} \sum_{i=1}^n \left\{ \left(\frac{r}{n} - \frac{r^2}{n^2} \right) T_1^2(\mathbf{z}_i) + \left(\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2} \right) \sum_{i' \neq i} T_1(\mathbf{z}_{i'}) T_1(\mathbf{z}_i) \right\}.
\end{aligned}$$

Again by utilizing the weak law of large number, we have

$$\frac{n}{r(n-r) \xi_{1,r}(\mathbf{x}_*)} \sum_{i=1}^n \left(\frac{r}{n} - \frac{r^2}{n^2} \right) T_1^2(\mathbf{z}_i) \xrightarrow{p} 1.$$

As for the term consisting of $T_1(\mathbf{z}_{i'})T_1(\mathbf{z}_i)$ above, we have

$$\begin{aligned} & \mathbb{E} \left\{ \frac{n}{r(n-r)\xi_{1,r}(\mathbf{x}_*)} \sum_{i=1}^n \left(\frac{r(r-1)}{n(n-1)} - \frac{r^2}{n^2} \right) \sum_{i' \neq i}^n T_1(\mathbf{z}_{i'}) T_1(\mathbf{z}_i) \right\}^2 \\ &= \mathbb{E} \left\{ \frac{1}{n(n-1)\xi_{1,r}(\mathbf{x}_*)} \sum_{i=1}^n \sum_{i' \neq i}^n T_1(\mathbf{z}_{i'}) T_1(\mathbf{z}_i) \right\}^2 \\ &= \frac{1}{n(n-1)} \rightarrow 0. \end{aligned}$$

Here, we use the fact $\mathbb{E}(\sum_{i' \neq i}^n T_1(\mathbf{z}_{i'}) T_1(\mathbf{z}_i))^2 = n(n-1)\xi_{1,r}^2(\mathbf{x}_*)$ which follows from a straightforward calculation. This implies that the interaction term converges to 0 in probability. With these two terms, we have

$$\frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n \frac{r(n-r)}{n^2} A_i \cdot T_1(\mathbf{z}_i) \xrightarrow{p} 1. \quad (\text{C.13})$$

Combing (C.11) and (C.12) with (C.13), we have

$$\frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n A_i^2 \xrightarrow{p} 1.$$

Moreover, it follows from (C.10) that

$$\mathbb{E} \frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n (M_i - A_i)^2 \rightarrow 0,$$

which means

$$\frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n (M_i - A_i)^2 \xrightarrow{p} 0.$$

The triangle inequality implies that

$$\begin{aligned} \frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n M_i^2 &\leq \frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \left(\sum_{i=1}^n A_i^2 + \sum_{i=1}^n (M_i - A_i)^2 \right) \xrightarrow{p} 1, \\ \frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n M_i^2 &\geq \frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \left(\sum_{i=1}^n A_i^2 - \sum_{i=1}^n (M_i - A_i)^2 \right) \xrightarrow{p} 1. \end{aligned}$$

Hence, the squeeze theorem gives that

$$\frac{n}{r^2\xi_{1,r}(\mathbf{x}_*)} \cdot \frac{n(n-1)}{(n-r)^2} \sum_{i=1}^n M_i^2 \xrightarrow{p} 1.$$

This completes the proof of Theorem 3.8.

D Proofs in Section C

We provide additional technical proofs for the conclusions in Section C. We first present the proof of Lemma C.4. Then we show several additional lemmas, which will be used to prove Proposition C.1.

Proof of Lemma C.4. By the definition we have

$$\psi(\eta) = \log \left(\int h(y) \exp(\eta y) dy \right),$$

and it is clear that

$$\mathbb{E}(Y|\eta) = \psi'(\eta), \quad \text{Var}(Y|\eta) = \psi''(\eta) \geq 0.$$

We conclude that $\psi(\cdot)$ is a convex function. By Assumption 3.1, the non-degeneracy and the existence of distinct values in the support of $p(y|\eta) = h(y) \exp(\eta y - \psi(\eta))$ ensure that $\text{Var}(Y|\eta) > 0$ for all $\eta \in [-C, C]$. This guarantees that the second derivative of the log-partition function, $\psi''(\eta) = \text{Var}(Y|\eta)$, is strictly positive. Furthermore, since $\mathbb{E}(Y|\eta)$ and $\mathbb{E}(Y^2|\eta)$ are defined as smooth integrals with respect to $p(y|\eta)$, they are continuous in η . As a result, $\psi''(\eta)$ is also continuous. By the compactness of $\eta \in [-C, C]$, there exists a constant $c > 0$ such that:

$$\psi''(\eta) \geq c > 0, \quad \text{for all } \eta \in [-C, C].$$

Using the second-order Taylor expansion and the definition of strong convexity, we establish:

$$\psi(x) - \psi(m) - \psi'(m)(x - m) = \frac{\psi''(\xi)}{2}(x - m)^2,$$

for some ξ between x and m . Since $C' \geq \psi''(\xi) \geq c' > 0$, we conclude:

$$\frac{C'}{2}(x - m)^2 \geq \psi(x) - \psi(m) - \psi'(m)(x - m) \geq \frac{c'}{2}(x - m)^2.$$

As for the Lipchitz property, by the equation $\psi'(\eta) = \mathbb{E}(Y|\eta)$, the bound of $\psi'(\eta)$ directly comes from the existence of the expectation of Y given η . \square

We propose several lemmas to address the general nonparametric regression cases, serving as the basis for the proof of Proposition C.1. Recall the definition of $R_n(\widehat{f}_n, f_0)$ in (2.5):

$$R_n(\widehat{f}_n, f_0) = \mathbb{E}[\ell(\mathbf{X}; \widehat{f}_n, f_0)],$$

and consider the construction of

$$\widehat{R}_n(\widehat{f}_n, f_0) = \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i; \widehat{f}_n, f_0)\right].$$

In the following lemmas, we demonstrate that the differences between $R_n(\widehat{f}_n, f_0)$ and $\widehat{R}_n(\widehat{f}_n, f_0)$ are small. Some of the proof techniques are inspired by Schmidt-Hieber (2020), but our analysis extends to a more general regression framework, and our proofs differ accordingly. In fact, the settings considered by Schmidt-Hieber (2020) are special cases within our broader context. Additionally, we provide upper and lower bounds for $\widehat{R}_n(\widehat{f}_n, f_0)$. These results will be used for proving Proposition C.1.

Lemma D.1. *Under Assumptions 3.1, it holds that*

$$\begin{aligned} |\widehat{R}_n(\widehat{f}_n, f_0) - R_n(\widehat{f}_n, f_0)| \leq & \frac{2R}{n} R_n^{1/2}(\widehat{f}_n, f_0) \sqrt{9n \log \mathcal{N}_n + 4n \cdot \frac{1 + 3 \log \mathcal{N}_n}{\log \mathcal{N}_n}} \\ & + \frac{2R(3 \log \mathcal{N}_n + 2)}{n} + 5\widetilde{\delta}C_{\text{Lip}}. \end{aligned}$$

Here, \mathcal{N}_n is the minimum number of $\widetilde{\delta}$ -covering of $\mathcal{F}(L, \mathbf{p}, s, \infty)$, and $R = 3FC_{\text{Lip}}$ is a constant assumed larger than 1 without loss of generality.

Lemma D.1 establishes the differences between $R_n(\widehat{f}_n, f_0)$ and $\widehat{R}_n(\widehat{f}_n, f_0)$. We now propose the following lemma, which addresses the challenges arising from the dependency

of ε_i on the covariates \mathbf{x}_i . This approach is different from previous work, such as [Schmidt-Hieber \(2020\)](#) and [Fan & Gu \(2024\)](#), which assumed independence between ε_i and \mathbf{x}_i .

Lemma D.2. *Suppose that Assumption 3.1 holds. Define $\varepsilon_i = y_i - \psi'(f_0(\mathbf{x}_i))$. For any estimator $\tilde{f} \in \mathcal{F}$, there exists constant $C > 0$ such that*

$$\left| \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \tilde{f}(\mathbf{x}_i) \right| \leq \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) + 2\tilde{\delta}C.$$

The main difficulty in Lemma D.2 is the dependence of ε_i and \mathbf{x}_i under the GNRM framework. Unlike traditional concentration inequalities, here we apply covering number analysis and truncation techniques to address the heteroskedasticity in GNRMs. Details can be referred to Section D.2. With Lemma D.2, we establish the following two lemmas, which provides the upper and lower bound of $\widehat{R}_n(\widehat{f}_n, f_0)$.

Lemma D.3. *Suppose that Assumption 3.1 holds. For any fixed $f \in \mathcal{F}$, it holds that*

$$\begin{aligned} \widehat{R}_n(\widehat{f}_n, f_0) \leq \inf_{f \in \mathcal{F}} \mathbb{E} \ell(\mathbf{X}; f, f_0) + \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) \\ + 2\tilde{\delta}C + \Delta_n(\widehat{f}_n). \end{aligned}$$

Here, \mathbf{X} is an independent copy of \mathbf{x}_i and $C > 0$ is a constant.

Lemma D.4. *Suppose that Assumption 3.1 holds. It holds that*

$$\widehat{R}_n(\widehat{f}_n, f_0) \geq \frac{1}{1+\rho} \cdot \left(\Delta_n(\widehat{f}_n) - 2\sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) - 4\tilde{\delta}C - \frac{C \log \mathcal{N}_n}{2n} \right)$$

for any $\rho > 0$. Here, $C > 0$ is a constant.

D.1 Proof of Lemma D.1

Given a minimum $\tilde{\delta}$ -covering of \mathcal{F} , let the center of the balls by f_j . Recall $\mathcal{D}_n = \{(\mathbf{x}_i, y_i), i \in [n]\}$. By construction there exists a j^* such that $\|\widehat{f} - f_{j^*}\|_\infty \leq \tilde{\delta}$. Moreover, by the assumptions we have $\|f_j\|_\infty \leq F$. Suppose that \mathbf{X}'_i ($i \in [n]$) are i.i.d random

variables with the same distribution of \mathbf{X} and independent of the sample $(\mathbf{x}_i)_{i \in [n]}$ (\mathcal{D}_n) and recall

$$\ell(\mathbf{x}; \widehat{f}_n, f_0) = -\psi'(f_0(\mathbf{x}))\widehat{f}_n(\mathbf{x}) + \psi(\widehat{f}_n(\mathbf{x})) + \psi'(f_0(\mathbf{x}))f_0(\mathbf{x}) - \psi(f_0(\mathbf{x})).$$

It is clear to see that $\ell(\mathbf{x}; \widehat{f}_n, f_0) \geq 0$ by Lemma C.4, and

$$\begin{aligned} & |R_n(\widehat{f}_n, f_0) - \widehat{R}_n(\widehat{f}_n, f_0)| \\ &= \left| \mathbb{E} \left\{ \frac{1}{n} \sum_{i=1}^n \left[\ell(\mathbf{X}'_i; \widehat{f}_n, f_0) - \ell(\mathbf{x}_i; \widehat{f}_n, f_0) \right] \right\} \right|. \end{aligned}$$

Rewriting the equation above by $\widehat{f}_n = \widehat{f}_n - f_{j^*} + f_{j^*}$, we have the inequality for any $\|\widehat{f}_n - f_{j^*}\|_\infty \leq \widetilde{\delta}$,

$$|\ell(\mathbf{x}; \widehat{f}_n, f_0) - \ell(\mathbf{x}; f_{j^*}, f_0)| \leq 2\widetilde{\delta}C_{\text{Lip}},$$

and hence

$$|R_n(\widehat{f}_n, f_0) - \widehat{R}_n(\widehat{f}_n, f_0)| \leq \mathbb{E} \left\{ \left| \frac{1}{n} \sum_{i=1}^n g_{j^*}(\mathbf{X}'_i, \mathbf{x}_i) \right| \right\} + 4\widetilde{\delta}C_{\text{Lip}}, \quad (\text{D.1})$$

where

$$g_{j^*}(\mathbf{X}'_i, \mathbf{x}_i) = \ell(\mathbf{X}'_i; f_{j^*}, f_0) - \ell(\mathbf{x}_i; f_{j^*}, f_0). \quad (\text{D.2})$$

We define g_j the same way with f_{j^*} replaced with f_j . By letting $U = \mathbb{E}^{1/2}\{\ell(\mathbf{X}; \widehat{f}_n, f_0)|\mathcal{D}_n\}$,

it is clear that

$$\begin{aligned} \left| \sum_{i=1}^n g_{j^*}(\mathbf{X}'_i, \mathbf{x}_i) \right| &= \frac{|\sum_{i=1}^n g_{j^*}(\mathbf{X}'_i, \mathbf{x}_i)|}{\max\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_{j^*}, f_0)|\mathcal{D}_n]\}} \cdot \max\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_{j^*}, f_0)|\mathcal{D}_n]\} \\ &\leq \max_j \frac{|\sum_{i=1}^n g_j(\mathbf{X}'_i, \mathbf{x}_i)|}{\max\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_j, f_0)]\}} \cdot \left(\sqrt{\log \mathcal{N}_n/n} + U + \widetilde{\delta}C_{\text{Lip}} \right), \end{aligned}$$

where the second inequality comes from the fact $\|f_{j^*} - \widehat{f}_n\|_\infty \leq \widetilde{\delta}$, and hence $\mathbb{E}^{1/2}[\ell(\mathbf{X}; f_{j^*}, f_0)|\mathcal{D}_n] \leq U + 2\widetilde{\delta}C_{\text{Lip}}$. Define the random variable

$$T = \max_j \frac{|\sum_{i=1}^n g_j(\mathbf{X}'_i, \mathbf{x}_i)|}{\max\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_j, f_0)]\}},$$

and it follows from (D.1) and the inequality above that

$$|R_n(\widehat{f}_n, f_0) - \widehat{R}_n(\widehat{f}_n, f_0)| \leq \mathbb{E} \frac{T}{n} \cdot \left(\sqrt{\log \mathcal{N}_n/n} + U + 2\widetilde{\delta}C_{\text{Lip}} \right) + 4\widetilde{\delta}C_{\text{Lip}}. \quad (\text{D.3})$$

Hence, it remains for us to consider the bound of (D.3). By the definition of $g_j(\mathbf{X}'_i, \mathbf{x}_i)$, it is clear that $\mathbb{E}g_j(\mathbf{X}'_i, \mathbf{x}_i) = 0$,

$$0 \leq \ell(\mathbf{x}; f_j, f_0) \leq R,$$

and hence $|g_j(\mathbf{X}'_i, \mathbf{x}_i)| \leq 2R$. Moreover, we have that

$$\begin{aligned} \text{Var}(g_j(\mathbf{X}'_i, \mathbf{x}_i)) &= 2\text{Var}(\ell(\mathbf{X}; f_j, f_0)) \\ &\leq 2\mathbb{E}\ell^2(\mathbf{X}; f_j, f_0) \leq 2R\mathbb{E}\ell(\mathbf{X}; f_j, f_0). \end{aligned}$$

As $g_j(\mathbf{X}'_i, \mathbf{x}_i)$ is independent with different i , then for any i.i.d bounded and centered random variables ξ_i with $|\xi_i| \leq M$, the Bernstein inequality shows that

$$\mathbb{P}\left(\left|\sum_{i=1}^n \xi_i\right| \geq t_0\right) \leq 2 \exp\left(-\frac{t_0^2}{2Mt_0/3 + 2\sum_{i=1}^n \text{Var}(\xi_i)}\right).$$

Combining the Bernstein inequality with the expression of T , we can apply the union bound and get that

$$\mathbb{P}(T/R \geq t) \leq 2\mathcal{N}_n \cdot \exp\left(-\frac{t^2}{\frac{4}{3}t \max^{-1}\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_j, f_0)]\} + 4n/R}\right).$$

Here, we take $\xi_i = g_j(\mathbf{X}'_i, \mathbf{x}_i)/R$, and let $t_0 = t \max\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_j, f_0)]\}$. According to (D.4), we can bound the value of $\mathbb{E}T$ and $\mathbb{E}T^2$, which will further be applied in (D.3). With the assumption above, we have $R \geq 1$, and the equation above can be simplified as

$$\mathbb{P}(T/R \geq t) \leq 2\mathcal{N}_n \cdot \exp\left(-\frac{t^2}{\frac{4}{3}t \max^{-1}\{\sqrt{\log \mathcal{N}_n/n}, \mathbb{E}^{1/2}[\ell(\mathbf{X}; f_j, f_0)]\} + 4n}\right). \quad (\text{D.4})$$

For the bound of $\mathbb{E}T/R$, it holds that

$$\begin{aligned}
\mathbb{E}T/R &= \mathbb{E}T/R \cdot \mathbf{1}\{T/R \leq 6\sqrt{n \log \mathcal{N}_n}\} + \mathbb{E}T/R \cdot \mathbf{1}\{T/R \geq 6\sqrt{n \log \mathcal{N}_n}\} \\
&= 6\sqrt{n \log \mathcal{N}_n} + \int_{6\sqrt{n \log \mathcal{N}_n}}^{+\infty} \mathbb{P}(T \geq t) dt \\
&\leq 6\sqrt{n \log \mathcal{N}_n} + \int_{6\sqrt{n \log \mathcal{N}_n}}^{+\infty} 2\mathcal{N}_n \exp\left(-\frac{t\sqrt{\log \mathcal{N}_n}}{2\sqrt{n}}\right) dt \\
&= 6\sqrt{n \log \mathcal{N}_n} + 4\sqrt{\frac{n}{\log \mathcal{N}_n}}. \tag{D.5}
\end{aligned}$$

Here, the inequality comes from (D.4) that when $t \geq 6\sqrt{n \log \mathcal{N}_n}$, $\mathbb{P}(T/R > t) \leq 2\mathcal{N}_n \exp\left(-\frac{t\sqrt{\log \mathcal{N}_n}}{2\sqrt{n}}\right)$, and the last inequality is by some algebra.

For the bound of $\mathbb{E}T^2/R^2$, it holds that

$$\begin{aligned}
\mathbb{E}T^2/R^2 &= \mathbb{E}T^2/R^2 \cdot \mathbf{1}\{T^2/R^2 \leq 6^2 n \log \mathcal{N}_n\} + \mathbb{E}T^2/R^2 \cdot \mathbf{1}\{T^2/R^2 \geq 6^2 n \log \mathcal{N}_n\} \\
&= 6^2 n \log \mathcal{N}_n + \int_{6^2 n \log \mathcal{N}_n}^{+\infty} \mathbb{P}(T \geq \sqrt{t}) dt \\
&\leq 6^2 n \log \mathcal{N}_n + \int_{6^2 n \log \mathcal{N}_n}^{+\infty} 2\mathcal{N}_n \exp\left(-\frac{\sqrt{t}\sqrt{\log \mathcal{N}_n}}{2\sqrt{n}}\right) dt \\
&= 36n \log \mathcal{N}_n + 16 \cdot \frac{1 + 3 \log \mathcal{N}_n}{\log \mathcal{N}_n}. \tag{D.6}
\end{aligned}$$

Here, the inequality comes from (D.4) that when $t \geq 6\sqrt{n \log \mathcal{N}_n}$, $\mathbb{P}(T/R > t) \leq 2\mathcal{N}_n \exp\left(-\frac{t\sqrt{\log \mathcal{N}_n}}{(4/3+4/(2R+2))\sqrt{n}}\right)$, and the last equality is by $\int_{b^2}^{+\infty} e^{-\sqrt{ta}} dt = 2(ab+1)e^{-ab}/a^2$.

Combining (D.5), (D.6) with (D.3), and noting that $\mathbb{E}U^2 = R_n(\widehat{f}_n, f_0)$ where U is defined below (D.2), we have

$$\begin{aligned}
|R_n(\widehat{f}_n, f_0) - \widehat{R}_n(\widehat{f}_n, f_0)| &\leq \mathbb{E}T \cdot \left(\sqrt{\log \mathcal{N}_n/n} + U + \widetilde{\delta}C_{\text{Lip}}\right) + 4\widetilde{\delta}C_{\text{Lip}} \\
&\leq \frac{1}{n} \mathbb{E}^{1/2} T^2 \cdot \mathbb{E}^{1/2} U^2 + \mathbb{E} \frac{T}{n} \cdot \left(\sqrt{\log \mathcal{N}_n/n} + \widetilde{\delta}C_{\text{Lip}}\right) + 4\widetilde{\delta}C_{\text{Lip}} \\
&\leq \frac{2R}{n} R_n^{1/2}(\widehat{f}_n, f_0) \sqrt{9n \log \mathcal{N}_n + 4n \cdot \frac{1 + 3 \log \mathcal{N}_n}{\log \mathcal{N}_n}} \\
&\quad + \frac{2R}{n} \left(3\sqrt{n \log \mathcal{N}_n} + 2\sqrt{\frac{n}{\log \mathcal{N}_n}}\right) \cdot \left(\sqrt{\log \mathcal{N}_n/n} + 2\widetilde{\delta}C_{\text{Lip}}\right) + 4\widetilde{\delta}C_{\text{Lip}} \\
&\leq \frac{2R}{n} R_n^{1/2}(\widehat{f}_n, f_0) \sqrt{9n \log \mathcal{N}_n + 4n \cdot \frac{1 + 3 \log \mathcal{N}_n}{\log \mathcal{N}_n}} + \frac{2R(3 \log \mathcal{N}_n + 2)}{n} + 5\widetilde{\delta}C_{\text{Lip}}.
\end{aligned}$$

Here, the first inequality is by the Cauchy-Schwarz inequality, the second inequality is by (D.5) and (D.6) and the last inequality holds when n is large enough.

D.2 Proof of Lemma D.2

For any estimator \tilde{f} taking values in \mathcal{F} , let j' be such that $\|\tilde{f} - f_{j'}\|_\infty \leq \tilde{\delta}$. We have

$$\begin{aligned} \left| \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \tilde{f}(\mathbf{x}_i) \right| &= \left| \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i (\tilde{f}(\mathbf{x}_i) - f_{j'}(\mathbf{x}_i) + f_{j'}(\mathbf{x}_i)) \right| \\ &\leq \frac{\tilde{\delta}}{n} \mathbb{E} \sum_{i=1}^n |\varepsilon_i| + \frac{1}{n} \left| \mathbb{E} \sum_{i=1}^n \varepsilon_i f_{j'}(\mathbf{x}_i) \right| \\ &\leq 2\tilde{\delta}C + \frac{1}{n} \left| \mathbb{E} \sum_{i=1}^n \varepsilon_i f_{j'}(\mathbf{x}_i) \right|. \end{aligned} \quad (\text{D.7})$$

Here, C is some large constant. the first inequality is by the triangle inequality, and the second inequality is by $\mathbb{E}|\varepsilon_i| = \mathbb{E}\{\mathbb{E}\{|y_i - \psi'(f_0(\mathbf{x}_i))|\}|\mathbf{x}_i\}\} \leq L + \mathbb{E}|Y| \leq C$ for some large constant C . It is remained for us to bound the last term above. To obtain the bound, we first introduce $r_j = \sqrt{\frac{\log \mathcal{N}_n}{n}} \vee \sqrt{\sum_{i=1}^n (f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2/n}$. Then the term can be written as

$$\begin{aligned} \frac{1}{n} \left| \mathbb{E} \sum_{i=1}^n \varepsilon_i f_{j'}(\mathbf{x}_i) \right| &= \frac{1}{n} \left| \mathbb{E} \sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i)) \right| \\ &\leq \frac{1}{n} \cdot \mathbb{E} \left[\left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + \sqrt{\sum_{i=1}^n (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2/n} \right) \cdot \frac{\left| \sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i)) \right|}{r_{j'}} \right] \\ &\leq \frac{1}{n} \cdot \mathbb{E} \left[\left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + \sqrt{\sum_{i=1}^n (\tilde{f}(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2/n + \tilde{\delta}} \right) \cdot \frac{\left| \sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i)) \right|}{r_{j'}} \right] \\ &\leq \frac{1}{n} \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + \mathbb{E}^{1/2} \left(\frac{\sum_{i=1}^n \ell(\mathbf{x}_i; \tilde{f}, f_0)}{cn} \right) + \tilde{\delta} \right) \cdot \mathbb{E}^{1/2} \frac{\left| \sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i)) \right|^2}{r_{j'}^2} \\ &= \frac{1}{n} \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + \hat{R}_n^{1/2}(\tilde{f}, f_0)/\sqrt{c} + \tilde{\delta} \right) \cdot \mathbb{E}^{1/2} \frac{\left| \sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i)) \right|^2}{r_{j'}^2}. \end{aligned} \quad (\text{D.8})$$

Here, the first equality is by $\mathbb{E}\varepsilon_i f_0(\mathbf{x}_i) = \mathbb{E}(\mathbb{E}\varepsilon_i f_0(\mathbf{x}_i)|\mathbf{x}_i) = 0$, the first inequality is by $r_j \leq \frac{\log \mathcal{N}_n}{n} + \sqrt{\sum_{i=1}^n (f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2/n}$, the second inequality is by triangle inequality $\|\mathbf{a} + \mathbf{b}\| \leq \|\mathbf{a}\| + \|\mathbf{b}\|$ and $\|f_{j'} - \tilde{f}\|_\infty \leq \tilde{\delta}$, and the last inequality is by Lemma C.4 that $\ell(\mathbf{x}; \tilde{f}, f_0) \geq c(\tilde{f}(\mathbf{x}) - f_0(\mathbf{x}))^2$ for some constant $c > 0$.

It remains for us to bound $\mathbb{E}^{1/2} \frac{|\sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i))|^2}{r_j^2}$. The heteroskedasticity in GNRMs also makes the bound of this term complicated, mainly due to the dependence among ε_i and \mathbf{x}_i . For simplicity in notation, define

$$\xi_j = \frac{|\sum_{i=1}^n \varepsilon_i (f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i))|}{r_j}. \quad (\text{D.9})$$

We investigate the noise behavior ξ_j^2 under different regimes to bound the term with heteroskedasticity. Specifically, we consider $\max_{j \in [\mathcal{N}_n]} \xi_j^2 \leq Cn \log \mathcal{N}_n$ and $\max_{j \in [\mathcal{N}_n]} \xi_j^2 \geq Cn \log \mathcal{N}_n$ conditional on \mathbf{x}_i . When $\max_{j \in [\mathcal{N}_n]} \xi_j^2 \leq Cn \log \mathcal{N}_n$ conditional on \mathbf{x}_i , the results will obviously hold. We mainly focus the regime when $\max_{j \in [\mathcal{N}_n]} \xi_j^2 \geq Cn \log \mathcal{N}_n$ conditional on \mathbf{x}_i . Conditional on $\mathbf{x}_1, \dots, \mathbf{x}_n$ [or $(\mathbf{x}_i)_i$ compactly], we have that ε_i 's ($i \in [n]$) are centered and independent sub-exponential random variables. Hence, we have

$$\begin{aligned} \text{Var}(\xi_j | (\mathbf{x}_i)_i) &= \frac{\sum_{i=1}^n \text{Var}(\varepsilon_i (f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i)) | (\mathbf{x}_i)_i)}{r_j^2} \\ &\leq \sum_{i=1}^n \mathbb{E}(\varepsilon_i^2 | (\mathbf{x}_i)_i) \cdot \mathbb{E}((f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2 | (\mathbf{x}_i)_i) / r_j^2 \\ &\leq 2n\kappa^2, \end{aligned}$$

where the first inequality is by $\text{Var}(U|X) \leq \mathbb{E}(U^2|X)$ for all random variables X and U , and the second inequality is by the definition of $r_j = \sqrt{\frac{\log \mathcal{N}_n}{n}} \vee \sqrt{\sum_{i=1}^n (f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i))^2/n}$, and $\mathbb{E}(\varepsilon_i^2 | (\mathbf{x}_i)_i) \leq 2\kappa^2$. Also, letting $a_{ij} = [f_j(\mathbf{x}_i) - f_0(\mathbf{x}_i)]/r_j$, we have

$$|a_{ij}| \leq 2F/r_j, \quad \sum_{i=1}^n a_{ij}^2 \leq n.$$

Therefore by the Bernstein inequality, it holds that

$$\mathbb{P}(|\xi_j| \geq t | (\mathbf{x}_i)_i) \leq 2 \exp \left\{ -c' \min \left\{ \frac{t^2}{n\kappa^2}, \frac{tr_j}{2F\kappa} \right\} \right\}$$

for some absolute constant $c' > 0$.

Let $\xi_{\max} = \max_{j \in [\mathcal{N}_n]} \xi_j$. By applying the union bound, it holds that

$$\begin{aligned} \mathbb{E}(|\xi_{\max}|^2 | (\mathbf{x}_i)_i) &= \int_0^{+\infty} \mathbb{P}(|\xi_{\max}|^2 \geq t | (\mathbf{x}_i)_i) dt \\ &\leq T + \int_T^{+\infty} \mathbb{P}(|\xi_{\max}| \geq \sqrt{t} | (\mathbf{x}_i)_i) dt \\ &\leq T + \mathcal{N}_n \cdot \int_T^{+\infty} 2 \exp \left\{ -c' \min \left\{ \frac{t}{n\kappa^2}, \frac{\sqrt{tr_j}}{2F\kappa} \right\} \right\} dt \\ &\leq T + \mathcal{N}_n \cdot \int_T^{+\infty} 2 \exp \left\{ -c' \min \left\{ \frac{t}{n\kappa^2}, \frac{\sqrt{t \log \mathcal{N}_n}}{2F\kappa\sqrt{n}} \right\} \right\} dt. \end{aligned}$$

Here, the last inequality is by $r_j \geq \sqrt{\frac{\log \mathcal{N}_n}{n}}$. By letting $T = Cn \log \mathcal{N}_n$ for some constant $C \geq 4F^2\kappa^2/c'^2$, it is clear that when $t \geq T$, $\frac{t}{n\kappa^2} \geq \frac{\sqrt{t \log \mathcal{N}_n}}{2F\kappa\sqrt{n}}$, the inequality above can be written as

$$\mathbb{E}(|\xi_{\max}|^2 | (\mathbf{x}_i)_i) \leq Cn \log \mathcal{N}_n + \mathcal{N}_n \cdot \int_{Cn \log \mathcal{N}_n}^{+\infty} 2 \exp \left\{ -c' \frac{\sqrt{t \log \mathcal{N}_n}}{2F\kappa\sqrt{n}} \right\} dt.$$

Again by utilizing the equality $\int_{b^2}^{+\infty} e^{-\sqrt{ta}} dt = 2(ab+1)e^{-ab}/a^2$ we conclude that as long as $C \geq 4F^2\kappa^2/c'^2$,

$$\begin{aligned} \mathbb{E}(|\xi_{\max}|^2 | (\mathbf{x}_i)_i) &\leq Cn \log \mathcal{N}_n + \mathcal{N}_n \cdot \int_{Cn \log \mathcal{N}_n}^{+\infty} 2 \exp \left\{ -c' \frac{\sqrt{t \log \mathcal{N}_n}}{2F\kappa\sqrt{n}} \right\} dt \\ &= Cn \log \mathcal{N}_n + 4\mathcal{N}_n \left(\frac{c'\sqrt{C}}{F\kappa} \log \mathcal{N}_n + 1 \right) e^{-\frac{c'\sqrt{C}}{2F\kappa} \log \mathcal{N}_n} \cdot \frac{F^2\kappa^2 n}{c'^2 \log \mathcal{N}_n} \\ &\leq 2Cn \log \mathcal{N}_n. \end{aligned}$$

By the definition of ξ_j in (D.9) and with $\xi_{\max} = \max_{j \in [\mathcal{N}]} \xi_j$, we have for any j' ,

$$\begin{aligned} \mathbb{E}^{1/2} \left| \frac{\sum_{i=1}^n \varepsilon_i (f_{j'}(\mathbf{x}_i) - f_0(\mathbf{x}_i))}{r_{j'}^2} \right|^2 &\leq \sqrt{\mathbb{E}(\mathbb{E}(|\xi_{\max}|^2 | (\mathbf{x}_i)_i))} \\ &\leq \sqrt{2Cn \log \mathcal{N}_n}. \end{aligned}$$

Combining the inequality above with (D.8), we have

$$\begin{aligned} \frac{1}{n} \left| \mathbb{E} \sum_{i=1}^n \varepsilon_i f_{j'}(\mathbf{x}_i) \right| &\leq \frac{1}{n} \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + \widehat{R}_n^{1/2}(\tilde{f}, f_0) / \sqrt{c} + \tilde{\delta} \right) \cdot \sqrt{2Cn \log \mathcal{N}_n}. \\ &\leq \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) \end{aligned} \quad (\text{D.10})$$

for some constant $C > 0$. Combining (D.10) with (D.7) we conclude that for any $\tilde{f} \in \mathcal{F}$,

$$\left| \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \tilde{f}(\mathbf{x}_i) \right| \leq \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) + 2\tilde{\delta}C$$

for some large constant $C > 0$, which completes the proof of Lemma D.2.

D.3 Proof of Lemma D.3

Recall the definition of $\Delta_n(\widehat{f}_n)$

$$\Delta_n(\widehat{f}_n) = \mathbb{E} \frac{1}{n} \sum_{i=1}^n -y_i \widehat{f}_n(\mathbf{x}_i) + \psi(\widehat{f}_n(\mathbf{x}_i)) - \inf_{f \in \mathcal{F}(L, \mathbf{p}, s, \mathcal{F})} \frac{1}{n} \sum_{i=1}^n -y_i f(\mathbf{x}_i) + \psi(f(\mathbf{x}_i)).$$

The first inequality we can obtain is, for any fixed $f \in \mathcal{F}$,

$$\mathbb{E} \frac{1}{n} \sum_{i=1}^n -y_i \widehat{f}_n(\mathbf{x}_i) + \psi(\widehat{f}_n(\mathbf{x}_i)) \leq \mathbb{E} \frac{1}{n} \sum_{i=1}^n -y_i f(\mathbf{x}_i) + \psi(f(\mathbf{x}_i)) + \Delta_n(\widehat{f}_n). \quad (\text{D.11})$$

Note that f and f_0 are both fixed. For any $\mathbf{X} \stackrel{\mathcal{D}}{=} \mathbf{x}_i$, the inequality above further implies that

$$\begin{aligned} \widehat{R}_n(\widehat{f}_n, f_0) &\leq \mathbb{E} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i; f, f_0) + \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \widehat{f}_n(\mathbf{x}_i) + \Delta_n(\widehat{f}_n) \\ &= \mathbb{E} \ell(\mathbf{X}; f, f_0) + \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \widehat{f}_n(\mathbf{x}_i) + \Delta_n(\widehat{f}_n) \\ &\leq \inf_{f \in \mathcal{F}} \mathbb{E} \ell(\mathbf{X}; f, f_0) + \sqrt{\frac{2C \widehat{R}_n(\widehat{f}_n, f_0) \log \mathcal{N}_n}{n}} \\ &\quad + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) + 2\tilde{\delta}C + \Delta_n(\widehat{f}_n). \end{aligned}$$

Here, $\varepsilon_i = y_i - \psi'(f_0(\mathbf{x}_i))$. The first inequality is from (D.11), the first equality is by $\mathbb{E} \ell(\mathbf{X}; f, f_0) = \mathbb{E} \ell(\mathbf{x}_i; f, f_0)$ for any fixed $f \in \mathcal{F}$, and the last inequality is by Lemma D.2.

D.4 Proof of Lemma D.4

Let $\tilde{f} = \operatorname{argmin}_{f \in \mathcal{F}} \sum_{i=1}^n -y_i f(\mathbf{x}_i) + \psi(f(\mathbf{x}_i))$ be the global empirical risk estimator. We have the following decomposition, by letting $\varepsilon_i = y_i - \psi'(f_0(\mathbf{x}_i))$. That is,

$$\widehat{R}_n(\widehat{f}_n, f_0) - \widehat{R}_n(\tilde{f}, f_0) = \Delta_n(\widehat{f}_n) + \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \widehat{f}_n(\mathbf{x}_i) - \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \tilde{f}(\mathbf{x}_i). \quad (\text{D.12})$$

From Lemma D.2, we have

$$\left| \mathbb{E} \frac{1}{n} \sum_{i=1}^n \varepsilon_i \tilde{f}(\mathbf{x}_i) \right| \leq \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) + 2\tilde{\delta}C$$

for any $\tilde{f} \in \mathcal{F}$ with some constant $C > 0$. Therefore (D.12) yields

$$\begin{aligned} \widehat{R}_n(\widehat{f}_n, f_0) - \widehat{R}_n(\tilde{f}, f_0) &\geq \Delta_n(\widehat{f}_n) - \sqrt{\frac{2C \widehat{R}_n(\widehat{f}_n, f_0) \log \mathcal{N}_n}{n}} - \sqrt{\frac{2C \widehat{R}_n(\tilde{f}, f_0) \log \mathcal{N}_n}{n}} \\ &\quad - 2\sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) - 4\tilde{\delta}C. \end{aligned} \quad (\text{D.13})$$

For any $a, b > 0$, it holds that $2ab \leq \rho a^2 + b^2/\rho$ for any $\rho > 0$. By letting $a = \widehat{R}_n(\widehat{f}_n, f_0)$, $b = \widehat{R}_n(\tilde{f}, f_0)$, $c = \sqrt{\frac{C \log \mathcal{N}_n}{2n}}$ and $d = \Delta_n(\widehat{f}_n) - 2\sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) - 4\tilde{\delta}C$, (D.13) can be written as

$$\begin{aligned} a - b &\geq d - 2\sqrt{ac} - 2\sqrt{bc} \\ &\geq d - \rho a - \frac{c^2}{\rho} - b - c^2 \end{aligned}$$

for any $\rho > 0$. We directly conclude that

$$a \geq \frac{d}{1 + \rho} - \frac{c^2}{1 + \rho},$$

which is

$$\widehat{R}_n(\widehat{f}_n, f_0) \geq \frac{1}{1 + \rho} \cdot \left(\Delta_n(\widehat{f}_n) - 2\sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n} \tilde{\delta}} \right) - 4\tilde{\delta}C - \frac{C \log \mathcal{N}_n}{2n} \right).$$

This completes the proof of Lemma D.4.

D.5 Proof of Proposition C.1

We first prove the upper bound. From Lemma D.1, we can let $a_1 = \widehat{R}_n(\widehat{f}_n, f_0)$, $b_1 = R_n(\widehat{f}_n, f_0)$, $c_1 = \frac{R}{n} \sqrt{9n \log \mathcal{N}_n + 4n \frac{1+3 \log \mathcal{N}_n}{\log \mathcal{N}_n}}$ and $d_1 = \frac{2R(3 \log \mathcal{N}_n + 2)}{n} + 5\widetilde{\delta}C_{\text{Lip}}$. The inequality in Lemma D.1 can be written as

$$|a_1 - b_1| \leq 2\sqrt{a_1}c_1 + d_1. \quad (\text{D.14})$$

Furthermore, from Lemma D.3, we can let $a_2 = a_1 = \widehat{R}_n(\widehat{f}_n, f_0)$, $c_2 = \sqrt{\frac{C \log \mathcal{N}_n}{2n}}$ and $d_2 = \inf_{f \in \mathcal{F}} \mathbb{E} \ell(\mathbf{X}; f, f_0) + \sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n}} \widetilde{\delta} \right) + 2\widetilde{\delta}C + \Delta_n(\widehat{f}_n)$. The inequality in Lemma D.3 can be written as

$$a_2 \leq 2\sqrt{a_2}c_2 + d_2. \quad (\text{D.15})$$

From (D.14), we have, for any $0 < \rho_1 < 1$,

$$b_1 - a_1 \leq \rho_1 a_1 + \frac{c_1^2}{\rho_1} + d_1,$$

while (D.15) implies that, for any $\rho_2 > 0$,

$$a_2 \leq \rho_2 a_2 + \frac{c_2^2}{\rho_2} + d_2.$$

Setting $a_1 = a_2$ and combining the two inequalities above, we have

$$b_1 \leq (1 + \rho_1) \cdot \left(\frac{c_2^2}{(1 - \rho_2)\rho_2} + d_2 \right) + \frac{c_1^2}{\rho_1} + d_1.$$

By the definition of a_1 , b_1 , c_1 , d_1 , a_2 , c_2 and d_2 and by letting $\rho_1 = 1$ and $\rho_2 = 0.5$, we conclude that there exists a constant C' depending only on F , function ψ and κ such that

$$R_n(\widehat{f}, f_0) \leq 2\Delta_n(\widehat{f}_n) + 2 \inf_{f \in \mathcal{F}} \mathbb{E} \ell(\mathbf{X}; f, f_0) + C' \left(\frac{\log \mathcal{N}_n}{n} + \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + 1 \right) \widetilde{\delta} \right).$$

We complete the upper bound in Proposition C.1.

We then prove the lower bound in Proposition C.1. From Lemma D.4, we can let $a_3 = \widehat{R}_n(\widehat{f}_n, f_0)$, and have

$$a_1 = a_3 \geq \frac{1}{1+\rho} \cdot \left(\Delta_n(\widehat{f}_n) - 2\sqrt{2C} \left(\frac{\log \mathcal{N}_n}{n} + \sqrt{\frac{\log \mathcal{N}_n}{n}} \widetilde{\delta} \right) - 4\widetilde{\delta}C - \frac{C \log \mathcal{N}_n}{2n} \right). \quad (\text{D.16})$$

Hence, from (D.14) we have for any $0 < \rho_3 < 1$,

$$a_1 - b_1 \leq \rho_3 a_1 + \frac{c_1^2}{\rho_3} + d_1,$$

or

$$b_1 \geq (1 - \rho_3) a_1 - \frac{c_1^2}{\rho_3} - d_1.$$

We can let $\rho = 0.2$ in (D.16) and $\rho_3 = 0.4$, and it is easy to conclude that there exists a constant c' depending only on F , function ψ and κ such that

$$R_n(\widehat{f}, f_0) \geq \frac{1}{2} \Delta_n(\widehat{f}_n) - c' \left(\frac{\log \mathcal{N}_n}{n} + \left(\sqrt{\frac{\log \mathcal{N}_n}{n}} + 1 \right) \widetilde{\delta} \right).$$