

DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS

ISSN 1440-771X

WORKING PAPER SERIES

**A Localised Neural network with
Dependent Data: Estimation and
Inference**

Jiti Gao, Bing Peng and Yanrong Yang

A Localized Neural Network with Dependent Data: Estimation and Inference

JITI GAO[†], BIN PENG[†] AND YANRONG YANG^{*}

[†]Monash University and ^{*}The Australian National University

June 9, 2023

Abstract

In this paper, we propose a localized neural network (LNN) model and then develop the LNN based estimation and inferential procedures for dependent data in both cases with quantitative/qualitative outcomes. We explore the use of identification restrictions from a nonparametric regression perspective, and establish an estimation theory for the LNN setting under a set of mild conditions. The asymptotic distributions are derived accordingly, and we show that LNN automatically eliminates the dependence of data when calculating the asymptotic variances. The finding is important, as one can easily use different types of wild bootstrap methods to obtain valid inference practically. In particular, for quantitative outcomes, the proposed LNN approach yields closed-form expressions for the estimates of some key estimators of interest. Last but not least, we examine our theoretical findings through extensive numerical studies.

Keywords: Activation Function; Binary Structure; Deep Learning; Kernel Regression; Time Series

JEL classification: C14, C22, C45

1 Introduction

Neural network (NN) architecture has received increasing attention over the last several decades. On relevant topics, a large number of papers have been published in different journals such as *Econometrica*, *The Annals of Statistics*, *Journal of Machine Learning Research*, *Neural Networks*, *Neurocomputing*, etc. by experts from different disciplines. Apparently, we cannot exhaust the literature, but refer the interested reader to Bartlett et al. (2021) and Fan et al. (2021) for extensive reviews from a methodological point of view.

NN usually includes three ingredients: input layer, hidden layer(s), and output layer. We now briefly comment on them one by one. The input layer is possibly the easiest one to understand, as it includes regressors only. Because NN usually concerns about a generic unknown function of multiple regressors of interest, it is unnecessary to involve any interaction terms among these regressors in the input layer. The hidden layer(s) include the parameters which need to be determined by training data. Once the parameters are estimated, one can load the test dataset in order to evaluate the performance of the newly constructed NN. Finally, the output layer receives the estimation/decision/prediction which varies under different context. Usually, there are two types of data for the output layer: quantitative outcome (e.g., y is continuous), and qualitative outcome (e.g., y is binary).

That said, in this paper we account for the features of input and output layers, and propose to consider two specific models as follows:

$$y = g(\mathbf{x}) + \varepsilon \quad \text{for continuous outcome,} \quad (1.1)$$

$$y = \begin{cases} 1 & g(\mathbf{x}) - \varepsilon \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{for binary outcome,} \quad (1.2)$$

in which $\mathbf{x} = (x_1, \dots, x_d)^\top$ includes d regressors with d being finite, and ε is the idiosyncratic error component. Further, we suppose that $g : [-a, a]^d \rightarrow \mathbb{R}$, where a is fixed¹. The main goal of this paper is to infer $g(\cdot)$ through a neural network architecture for both models. In addition, for Model (1.2) only, we suppose that the respective probability density function (PDF) and the cumulative distribution function (CDF) of ε are known, and denote the PDF and CDF by $\phi_\varepsilon(\cdot)$ and $\Phi_\varepsilon(\cdot)$ respectively. Here, the information about $\phi_\varepsilon(\cdot)$ and $\Phi_\varepsilon(\cdot)$ is necessary for carrying on likelihood estimation. Possibly, one may further merge both models into one framework (say, $y = g(\mathbf{x}, \varepsilon)$), and consider a generic loss function, such as Chen and White (1999) and Section 1 of Bartlett et al. (2021), which however will involve some high level conditions to facilitate the development. As we will derive some results which are for Model (1.2) in particular and

¹We work with fixed a in the main text in order to be consistent with the majority of the existing NN literature, and explain how to allow $g(\cdot)$ to be defined on \mathbb{R}^d in Appendix A.3.

will also show that Model (1.1) may enjoy some closed-form estimation procedures, we do not further pursue a more general setup in this paper, but would like to investigate the general setup in future research.

Until very recently, the investigation on NN architecture mainly focuses on some fixed design (e.g., Cybenko, 1989 and many follow-up studies since then), or uses independent and identically distributed (i.i.d.) data (e.g., Kohler and Krzyzák (2017), Bauer and Kohler, 2019, Schmidt-Hieber, 2020, and many references therein). There are only limited studies available for us to understand NN with dependent data from theoretical perspective (see, Chen and Shen, 1998, Chen, 2007, for example), although NN based methods have been widely used to study different types of time series data in practice (e.g., Hill et al., 1996, Chen et al., 2001, Gu et al., 2021, Gu et al., 2020, just to name a few). We would like to contribute along this line of research, so throughout we assume the following time series data are observable:

$$\{(y_t, \mathbf{x}_t) \mid t \in [T]\}, \quad (1.3)$$

where, for a positive integer T , $[T]$ stands for $\{1, \dots, T\}$. The values of $\{y_t\}$ vary with respect to the model being studied. We shall not mention it again unless misunderstanding may arise.

Up to this point, it is worth further commenting on the hidden layer, which usually involves lots of neurons. Here, a neuron stands for an activation function, which maps a linear combination of the regressors from the input layer to a certain range of the real line. As a consequence, lots of parameters are actually involved in the hidden layer. The literature seems to agree that in order to ensure the number of effective parameters within a reasonable range, sparsity should be imposed (e.g., Schmidt-Hieber, 2020; Wang and Lin, 2021; Fan and Gu, 2022; Zhong et al., 2022, just to name a few), and in particular Schmidt-Hieber (2020) writes “*the network sparsity assumes that there are only few non-zero/active network parameters*”, which essentially is reflected by the number of active neurons. In this paper we consider Sigmoidal function, $\sigma(\cdot) : \mathbb{R} \rightarrow [0, 1]$, as the activation function, and acknowledge the growing literature on the non-sigmoidal rectified linear unit (ReLU) function. We conjecture using ReLU function one may derive results similar to what we are about to present below, but it requires some careful investigation in future research.

All things considered, we draw Figure 1 for the purpose of illustration, which presents a neural network architecture with only a small number of active neurons. Specifically, the elements of \mathbf{x} form the input layer, while, in the hidden layer, only neurons in the dark area are activated. Eventually, a final result is passed on to the output layer. To understand Figure 1 fully, a few questions arise naturally:

1. Provided a set of dependent time series data, how do we train a neural network, such as that in Figure 1 ?

2. Why are there only a small number of neurons getting activated ?
3. Given a set of activated neurons, can we further reduce the number of parameters ? In other words, how does sparsity come to play just among the activated neurons ?
4. When the least absolute shrinkage and selection operator (i.e., LASSO) is employed, how do we define the set of true parameters ?
5. Can any inference (such as a confidence interval) be established ? and so forth.

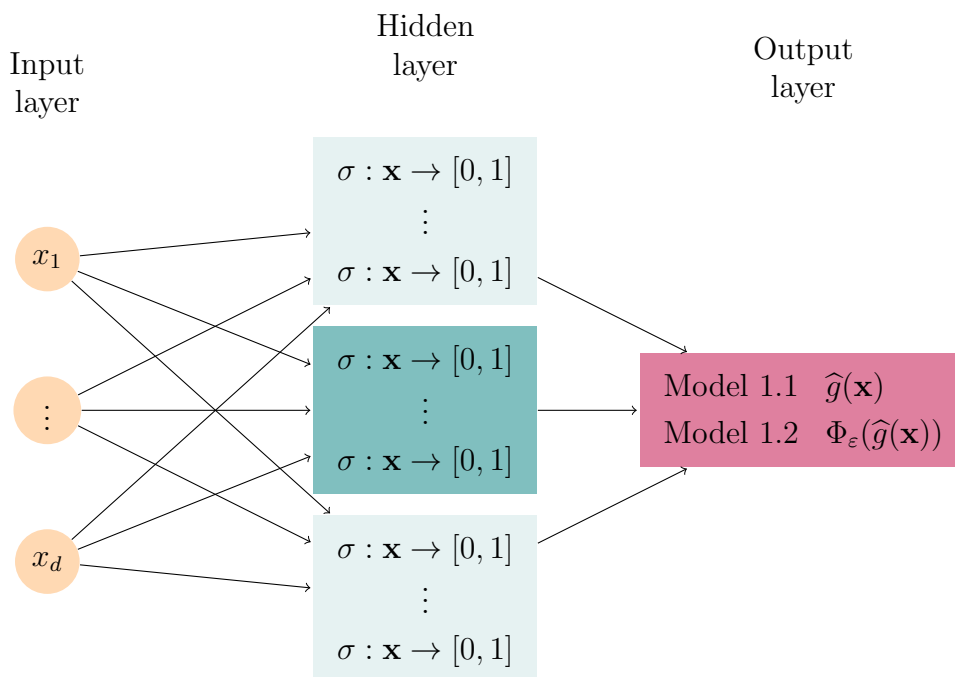


Figure 1: A Neural Network with Sparsity

In Figure 1, $\hat{g}(\cdot)$ is an estimated version of $g(\cdot)$.

With respect to the existing literature, it seems that so far the entire literature has been focusing on prediction errors, and barely talks about how to build feasible inferential procedures, such as constructing confidence intervals. A few exceptions known to us are Du et al. (2021), Farrell et al. (2021) and Chen et al. (2022) for example on estimation and inference for the average treatment effect rather than on $g(\cdot)$ we are interested in this paper. In other words, to the best of our knowledge, there have been neither asymptotic distribution results available for $\hat{g}(\cdot)$ nor inferential results available for $g(\cdot)$. Therefore, the main objective and contribution of this paper is that we develop the LNN based approach to addressing important estimation and inferential issues for $g(\cdot)$ associated with time series data.

Another challenge which arises with the complexity of the NN architecture is the transparency of algorithms. It is actually associated with our discussion about lacking of understanding on how to establish a feasible version practically. Although a variety of software

packages (e.g., Günther and Fritsch, 2010 and references therein) have been well adopted, the detailed implementation is still very vague in our view. When writing up this paper, we cannot find any existing paper using NN approach on social science topics that provides description about numerical implementation. While we agree with Athey (2019) that “*machine learning approaches will have dramatic impacts on different fields of social science within a short time frame*”, a transparent algorithm is much needed to ensure that findings are relevant and useful in practice.

Below, we provide some examples to illustrate our concern about the transparency of algorithms further. In the literature, the “`neuralnet`” R package (see Günther and Fritsch, 2010 for detailed illustration) has been well adopted. When training a NN, a key parameter is called “`hidden`” that is *a vector of integers specifying the number of hidden neurons in each layer* by R document, and the package refers to Murata et al. (1994) regarding the choice of the number of neurons. However, it is worth pointing out that Murata et al. (1994) use a modified AIC criterion to investigate the case with the number of neurons (as well as the number of parameters) being finite which is reflected in their asymptotic development given in the appendix. As a consequence, the arguments made in Murata et al. (1994) no longer hold when the number of neurons is diverging. As we are about to show below, having a diverging number of neurons is the minimum requirement to achieve asymptotic consistency, and the rate of divergence is also associated with the sample size under a set of minor conditions. Similar issues also apply to “`deepnet`” package of R, “`torch.nn.Linear`” and “`tfl.layers.Linear`” of Python, “`feedforwardnet`” of Matlab, etc. All in all, the fundamental question behind these is how to establish NN based estimation and inference for practical implementation.

Having those said, our contributions are four-fold. First, we explore the use of identification restrictions from a nonparametric regression perspective, and establish the LNN based estimation method for the effective parameters under a set of mild conditions. Second, asymptotic distributions are derived accordingly for inferential purposes, and we show that the LNN method automatically eliminates the dependence of data when calculating the asymptotic variances. The finding is important, as one can easily use different types of wild bootstrap methods to obtain valid inference practically. Also, it is worth emphasizing that for the first time in the relevant literature, we are able to rigorously establish asymptotic distributions for the proposed estimators of the unknown functions using the LNN approach. Third, for quantitative outcomes, the proposed LNN approach yields closed-form expressions for the estimators of the parameters of interest. Last but not least, we examine our theoretical findings through extensive numerical studies.

The rest of paper is organised as follows. Section 2 provides some basic assumptions with justifications. Section 3 first introduces the LNN architecture from a nonparametric regression

perspective in Section 3.1, and then establishes asymptomatic results associated with Models (1.1) and (1.2) in Section 3.2 and Section 3.3 respectively. Section 3.4 summarizes some key points outlined in this paper, and further discusses some important issues associated with LNN. We provide extensive simulation studies in Section 4 to examine the theoretical findings. Section 5 includes two empirical studies using Model (1.1) and Model (1.2) respectively. Section 6 concludes. We provide some additional information about Sigmoidal squasher in Appendix A.1, some extra simulation results in Appendix A.2, an important extension from $\mathbf{x} \in [0, 1]^d$ to $\mathbf{x} \in \mathbb{R}^d$ in Appendix A.3, important lemmas in Appendix A.4, and all the necessary proofs in Appendix A.5.

2 Notation and Assumptions

In this section, we introduce some notation that will be repeatedly used throughout the rest of this paper, and lay out some necessary assumptions.

We start from mathematical symbols. Vectors and matrices are always expressed in bold font. Further, $I(\cdot)$ defines the indicator function; $\|\cdot\|$ denotes the Euclidean norm of a vector or the Frobenius norm of a matrix; $\mathbf{0}_a$ and $\mathbf{1}_a$ are respectively $a \times 1$ vectors of zeros and ones for $a \in \mathbb{N}$; for a vector of nonnegative integers $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d)^\top \in \mathbb{N}_0^d$ in which $\mathbb{N}_0 = 0 \cup \mathbb{N}$, let $\boldsymbol{\alpha}! = \alpha_1! \cdots \alpha_d!$; \mathbf{c} , \mathbf{C} and $O(1)$ always stand for fixed constants, and may be different at each appearance; \rightarrow_P and \rightarrow_D stand for convergence in probability and convergence in distribution, respectively. For $g(\mathbf{x})$ defined in Model (1.1) and Model (1.2), if the partial derivative of $g(\mathbf{x})$ exists, we write $\frac{\partial^{|\boldsymbol{\alpha}|} g(\mathbf{x})}{\partial \mathbf{x}^{\boldsymbol{\alpha}}} = \frac{\partial^{|\boldsymbol{\alpha}|} g(\mathbf{x})}{\partial x_1^{\alpha_1} \cdots \partial x_d^{\alpha_d}}$ for short, where $|\boldsymbol{\alpha}| = \sum_{j=1}^d \alpha_j$. We let $\|g\|_\infty = \sup_{\mathbf{x} \in [-a, a]^d} |g(\mathbf{x})|$.

Having these symbols in hand, we are ready to proceed. First, we regulate the function of interest and the activation function. Below, we present a definition, and then formally state the first assumption of the paper.

Definition 2.1 (Continuity). *Let $p = q + s$ for some $q \in \mathbb{N}$ and $0 < s \leq 1$. A function $m : [-a, a]^d \mapsto \mathbb{R}$ is called (p, \mathcal{C}) -smooth, if for every $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $|\boldsymbol{\alpha}| = q$ the partial derivative $\frac{\partial^{|\boldsymbol{\alpha}|} m(\mathbf{x})}{\partial \mathbf{x}^{\boldsymbol{\alpha}}}$ exists and satisfies that for all $\mathbf{x}, \mathbf{z} \in [-a, a]^d$,*

$$\left| \frac{\partial^{|\boldsymbol{\alpha}|} m(\mathbf{x})}{\partial \mathbf{x}^{\boldsymbol{\alpha}}} - \frac{\partial^{|\boldsymbol{\alpha}|} m(\mathbf{z})}{\partial \mathbf{z}^{\boldsymbol{\alpha}}} \right| \leq \mathcal{C} \cdot \|\mathbf{x} - \mathbf{z}\|^s.$$

Definition 2.1 basically defines a family of sufficiently smooth functions.

Assumption 1.

1. Let $g(\mathbf{x})$ be (p, \mathcal{C}) -smooth as in Definition 2.1, and $\max_{|\boldsymbol{\alpha}| \leq q} \left\| \frac{\partial^{|\boldsymbol{\alpha}|} g(\mathbf{x})}{\partial \mathbf{x}^{\boldsymbol{\alpha}}} \right\|_\infty \leq \mathbf{c}$.

2. Let Sigmoidal function $\sigma(\cdot) : \mathbb{R} \rightarrow [0, 1]$ satisfy that

(a) $\sigma(\cdot)$ is at least $q + 1$ times continuously differentiable with bounded derivatives.

(b) A point $u_\sigma \in \mathbb{R}$ exists, where all derivatives up to the order q of $\sigma(\cdot)$ are different from zero.

Assumption 1.1 is widely adopted in the literature of nonparametric regression (e.g., Li and Racine, 2007). The main point is that each component in Taylor expansion of $g(\cdot)$ is bounded and also sufficiently smooth.

Assumption 1.2 imposes very limited conditions on the activation function, and nests a wide class of activation functions commonly used in the relevant literature as special cases, e.g., Sigmoidal squasher (i.e., $\sigma(x) = \frac{1}{1+\exp(-x)}$), Error function (i.e., $\sigma(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-w^2)dw$), etc. We refer the interested reader to Dubey et al. (2022) for a comprehensive review on different activation functions. In addition, we note that the activation function is usually chosen by users, so in general it can be much smoother than the function of interest (i.e., $g(\mathbf{x})$). A detailed example is provided in the online supplementary Appendix A.1 to show Assumption 1.2 can be easily fulfilled in practice. Obviously, Assumption 1.2 rules out ReLU function, which will be left for future research as explained in the introduction.

We then regulate the behaviour of the time series involved in (1.3).

Assumption 2.

1. $\{(\mathbf{x}_t, \varepsilon_t) \mid t \in [T]\}$ are strictly stationary and α -mixing with mixing coefficient

$$\alpha(t) = \sup_{A \in \mathcal{F}_{-\infty}^0, B \in \mathcal{F}_t^\infty} |P(A)P(B) - P(AB)|$$

satisfies $\sum_{t=1}^\infty \alpha^{\nu/(2+\nu)}(t) < \infty$ for some $\nu > 0$, which is the same as involved in Assumption 2.3 below, where $\mathcal{F}_{-\infty}^0$ and \mathcal{F}_t^∞ are the σ -algebras generated by $\{(\mathbf{x}_s, \varepsilon_s) : s \leq 0\}$ and $\{(\mathbf{x}_s, \varepsilon_s) : s \geq t\}$, respectively.

2. The density function $f_{\mathbf{x}}(\cdot)$ of \mathbf{x}_1 is Lipschitz continuous on $[-a, a]^d$, and is also bounded away from 0 below and bounded away from ∞ above on $[-a, a]^d$.

3. Model (1.1): $E[\varepsilon_1 \mid \mathbf{x}_1] = 0$, $E[\varepsilon_1^2 \mid \mathbf{x}_1] = \sigma_\varepsilon^2$ almost surely (a.s.), and $E[|\varepsilon_1|^{2+\nu} \mid \mathbf{x}_1] \leq \mathbf{c}$ a.s.

Model (1.2): $\phi_\varepsilon(\cdot)$ and $\Phi_\varepsilon(\cdot)$ are known, and $[1 - \Phi_\varepsilon(g(\mathbf{x}))]\Phi_\varepsilon(g(\mathbf{x})) \neq 0$ uniformly on $[-a, a]^d$.

Assumption 2 is standard in the literature of time series regression (Fan and Yao, 2003; Gao, 2007). As we shall embed the regressors $\{\mathbf{x}_t\}$ in the activation function under consideration

that is bounded from both below and above, we do not impose many assumptions on the behaviour of $\{\mathbf{x}_t\}$. In a sense, LNN automatically rescales the regressors by design, which somewhat simplifies our asymptotic analysis. Also, we note that the uncorrelated structure imposed between $\{\mathbf{x}_t\}$ and $\{\varepsilon_t\}$ is not necessary in this paper. For Model (1.1), we in fact can allow for heterogeneity, such as

$$E[\varepsilon_1^2 | \mathbf{x}_1] = \psi(\mathbf{x}_1),$$

which however makes notation even more complicated than what we need to involve. Similarly, for Model (1.2), we may also introduce heterogeneity in a more general setup such as

$$\Pr(y_1 = 1 | \mathbf{x}_1) = \Phi_\varepsilon(g(\mathbf{x}_1) | \tilde{\psi}),$$

in which $\tilde{\psi}$ stands for an unknown variance of ε_1 to be determined by data. Again, it will complex our discussion and deviate from our main goal.

Finally, we would like to point out that $\{\mathbf{x}_t\}$ do not have to be strictly stationary, and can have other more complex structures such as linear processes, locally stationarity, heterogeneity, deterministic trends, etc. Surely, the corresponding development and asymptotic results will need to be modified accordingly, but it will not add too many extra credits to the original idea of this paper. Therefore, we adopt the current setting throughout the rest of this paper.

We now move on to construct the so-called LNN architecture for both models, and will provide some necessary comments along the way.

3 Methodology and Theory

In this section, we first introduce the LNN architecture in Section 3.1, and then establish asymptotic results for both models in Sections 3.2 and 3.3 respectively. Section 3.4 further discusses some issues which may guide our future research.

One important goal of this paper is to understand the logic of NN. Therefore, it is worth to be more specific on how NN architecture works conceptually. In the literature of machine learning, one prefers to target the entire area that the unknown function is defined on, and normally uses training set to pre-specify a lot of parameters, which do not change with respect to different observations in the test set. By doing so, one just needs to pay some price when calculating the parameters in the first time, and no longer needs to update them with respect to different values in the test set. Having this mind, we are ready to consider the estimation of both Models (1.1) and (1.2).

3.1 The Localized Neural Network Architecture

In contrast to the current literature that usually allows for all parameters to be estimated from data, we start by presenting some identification conditions, which can help reduce the number of effective parameters significantly. In a sense, we bring in sparsity from a different perspective.

In view of the literature (e.g., Bauer and Kohler, 2019; Schmidt-Hieber, 2020), one key idea behind NN architecture is that it can approximate polynomial terms, of which as well understood the linear combination can further approximate unknown functions by standard nonparametric analysis. In this paper, we do the same, but the difference is that we introduce a bandwidth parameter h below, which serves the same purpose as that employed in nonparametric kernel estimation. This is basically why we define the neural network under consideration as a localized neural network (LNN). More importantly, the length of h has a direct control on activated and non-activated neurons involved in the hidden layer. We shall be clear on this point when presenting the estimation method in Sections 3.2 and 3.3.

To proceed further, we develop several new and useful results in Lemmas 3.1-3.3 to show how to approximate $g(\cdot)$ by a localized neural network system.

Lemma 3.1 (Identification). *Suppose that Assumption 1.2 holds. For $\forall x_0 \in \mathbb{R}^d$, we can find $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{q+1})^\top$ and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{q+1})^\top$ such that*

$$\sup_{|x-x_0| \leq h} \left| \sum_{k=1}^{q+1} \gamma_k \sigma(\beta_k \cdot (x - x_0) + u_\sigma) - (x - x_0)^q \right| = O(h^{q+1})$$

holds, where \mathbf{C} is a constant, h is a bandwidth, and $\gamma_k = \frac{(-1)^{q+k-1} \mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \binom{q}{k-1}$ and $\beta_k = \frac{k-1}{\mathbf{C}}$ for $k \in [q+1]$.

Remark 3.1.

1. Note that Lemma 3.1 is independent of data. Only the number of polynomial terms to be approximated depends on the smoothness of $g(\cdot)$. In fact, even the smoothness is predetermined by users, which is equivalent to choosing among local constant, local linear, and local higher order polynomials in the literature of kernel regression (Fan and Gijbels, 1996). That said, $\boldsymbol{\gamma}$, $\boldsymbol{\beta}$ and u_σ are fully decided by the activation function and the polynomial terms, so they are known prior to regression.

The constant \mathbf{C} raises an issue of identifiability, so we simply let $\mathbf{C} = 1$ throughout. Also, as required by Assumption 1.2, $\sigma^{(q)}(u_\sigma) \neq 0$, which can be easily realized in view of Figure A.1 of the online supplementary appendix.

2. h is equivalent to the bandwidth of kernel regression, so it will be decided by the sample size eventually. Later on, we show that h reflects the relationship between the sample size and the total number of neurons, and also controls how many neurons are activated.
3. $k - 1$ of β_k is to invoke the property of the Stirling number of the second kind in the development. As a result, the LNN architecture with a suitable Sigmoidal function can approximate the polynomial terms, of which the linear combination further ensures an unknown smooth function can be approximated.
4. Lemma 3.1 in fact yields a recursive relationship, as one can repeatedly invoke Lemma 3.1 to replace $(x - x_0)$ inside the activation. It will then yield a LNN architecture with multiple hidden layers. However, we do not see any benefit of doing so unless $g(\cdot)$ in both models has certain specific structure. That said, we focus on LNN with one hidden layer in this study.

To carry on, we further introduce a few more symbols. Denote that

$$\mathcal{P}_q = \left\{ \text{Linear span of the monomials } \prod_{k=1}^d x_k^{n_k} \text{ with } 0 \leq |\mathbf{n}| \leq q \right\}, \quad (3.1)$$

where $\mathbf{n} = (n_1, \dots, n_d)^\top \in \mathbb{N}_0^d$ and $|\mathbf{n}| = \sum_{i=1}^d n_i$. The dimension of \mathcal{P}_q is equivalent to the number of distinctive terms in the expansion of $(1 + x_1 + \dots + x_d)^q$, so

$$\dim \mathcal{P}_q = \binom{d+q}{d} := d_q,$$

where d_q is introduced for notational simplicity, and is fixed due to the fact that both d and q are finite. Furthermore, we define

$$\mathbf{m}(\mathbf{x} | \mathbf{x}_0) = (m_1(\mathbf{x} | \mathbf{x}_0), \dots, m_{d_q}(\mathbf{x} | \mathbf{x}_0))^\top, \quad (3.2)$$

where $m_j(\mathbf{x} | \mathbf{x}_0)$'s are the basis monomials (centred at \mathbf{x}_0) in \mathcal{P}_q . Also, we denote a set $C_{\mathbf{x}_0, h}$:

$$C_{\mathbf{x}_0, h} = \{\mathbf{x} \mid |x_j - x_{0,j}| \leq h \text{ for } j \in [d]\}, \quad (3.3)$$

where x_j and $x_{0,j}$ are the j^{th} elements of \mathbf{x} and \mathbf{x}_0 respectively.

With these symbols in hand, we present the next lemma, which demonstrates the feasibility of the LNN architecture.

Lemma 3.2 (Feasibility). *Suppose that Assumption 1.2 holds. For $\forall \mathbf{x}_0 \in \mathbb{R}^d$, we define $p(\mathbf{x} | \mathbf{x}_0, \boldsymbol{\lambda}) = \boldsymbol{\lambda}^\top \mathbf{m}(\mathbf{x} | \mathbf{x}_0)$, where $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_{d_q})^\top$ with $\|\boldsymbol{\lambda}\| \leq c$. Then a LNN of the type*

$$s(\mathbf{x} | \mathbf{x}_0, \tilde{\boldsymbol{\lambda}}) = (\tilde{\boldsymbol{\lambda}} \otimes \boldsymbol{\gamma})^\top \boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_0) \quad \text{with} \quad \boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_0) = \{\sigma([1, \mathbf{x}^\top - \mathbf{x}_0^\top] \boldsymbol{\pi}_j)\}_{d_q(q+1) \times 1}$$

exists such that

$$\sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} |s(\mathbf{x} | \mathbf{x}_0, \tilde{\boldsymbol{\lambda}}) - p(\mathbf{x} | \mathbf{x}_0, \boldsymbol{\lambda})| = O(h^{q+1}),$$

where $\tilde{\boldsymbol{\lambda}} = \mathbf{D}^\top \boldsymbol{\lambda}$ with \mathbf{D} being a rotation matrix, and

$$(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_{d_q(q+1)}) = \frac{\text{diag}\{h, \mathbf{I}_d\} \mathbf{W} (\mathbf{I}_{d_q} \otimes \boldsymbol{\beta}^\top)}{d+1} + \begin{bmatrix} u_\sigma \\ \mathbf{0}_d \end{bmatrix} \otimes \mathbf{1}_{d_q(q+1)}^\top,$$

in which \mathbf{W} is a user chosen matrix satisfying that $\max_j \|\mathbf{w}_j\| \leq \sqrt{d+1}$ and $\mathbf{w}_{j_1} \neq \mathbf{w}_{j_2}$ for any given $j_1, j_2 \in [d_q]$, and \mathbf{w}_j stands for the j^{th} column of \mathbf{W} .

Remark 3.2.

1. Note that we use vector form to rewrite each element of $\boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_0)$, i.e.,

$$\sigma([1, \mathbf{x}^\top - \mathbf{x}_0^\top] \boldsymbol{\pi}_j)$$

which is exactly what Sigmoidal activation function does (i.e., mapping a linear combination of regressors plus a location parameter to $[0, 1]$). By design, $\boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_0)$ naturally and automatically explores different interaction terms of the regressors.

2. By choosing $\boldsymbol{\lambda}$ and \mathbf{x}_0 , $p(\mathbf{x} | \mathbf{x}_0, \boldsymbol{\lambda})$ can approximate $g(\mathbf{x})$ reasonably well in a small neighbourhood of \mathbf{x}_0 . This is not hard to see, as the leading terms of the q^{th} order Taylor expansion of $g(\mathbf{x})$ at \mathbf{x}_0 can be written as follows:

$$g(\mathbf{x}) \simeq \sum_{0 \leq |\mathbf{J}| \leq q} \frac{1}{\mathbf{J}!} \cdot \frac{\partial^{|\mathbf{J}|} g(\mathbf{x}_0)}{\partial \mathbf{x}^{\mathbf{J}}} (\mathbf{x} - \mathbf{x}_0)^{\mathbf{J}} := \boldsymbol{\lambda}^\top \mathbf{m}(\mathbf{x} | \mathbf{x}_0) := p(\mathbf{x} | \mathbf{x}_0, \boldsymbol{\lambda}),$$

where the definition of $\boldsymbol{\lambda}$ should be obvious in view of the definition of $\mathbf{m}(\mathbf{x} | \mathbf{x}_0)$ according to (3.2). As a result, $s(\mathbf{x} | \mathbf{x}_0, \tilde{\boldsymbol{\lambda}})$ essentially approximates $g(\mathbf{x})$ but in a small neighbourhood only, and the rate produced in Lemma 3.2 is proportional to h^{q+1} as a bias term.

3. Lemma 3.2 infers that the LNN architecture rotates the parameters of interest (i.e., $\boldsymbol{\lambda}$) with a pre-determined $d_q \times d_q$ full rank matrix \mathbf{D} , so the new parameters of interest become $\tilde{\boldsymbol{\lambda}}$. Therefore, accounting for identification conditions properly ensures the number of effective parameters to be estimated is far less than it looks like. Without addressing identification issues, one has to determine $d_q(q+1)(d+2)$ parameters. In other words, utilizing identification conditions automatically reduces the number of effective parameters which need to be estimated.

4. Without loss of generality, we can let $\mathbf{w}_j = \frac{\sqrt{d+1}}{q} \mathbf{w}_j^*$, where $\{\mathbf{w}_j^*\}$ are the vectors corresponding to the powers of the distinctive terms in the expansion of $(1 + x_1 + \cdots + x_d)^q$. Thus, $\|\mathbf{w}_j\| \leq \frac{\sqrt{d+1}}{q} |\mathbf{w}_j^*| = \sqrt{d+1}$, which further ensures Lemma 3.1 can be invoked.

Using Lemma 3.2, the final form of LNN requires us to focus on a compact subspace of \mathbb{R}^d (i.e., $[-a, a]^d$), as we need to partition it into lots of small cubes. As discussed in Appendix A.3 below, the main results remain valid when allowing $a \rightarrow \infty$ along with the sample size. These small cubes may have different names in each appearance. For example, they are referred to as (hyper-)cubes in Bauer and Kohler (2019) and Schmidt-Hieber (2020), and are referred to as localization in Farrell et al. (2021). In our view, each cube is corresponding to an effective sample set, so it is very much similar to the effective sample range in nonparametric kernel regression, which is jointly determined by the point of interest and the bandwidth.

That said, for a given integer $M \geq 1$, we subdivide $[-a, a]^d$ into M^d cubes of side length as follows:

$$h = \frac{a}{M},$$

and for comprehensibility, we number these cubes by $C_{\mathbf{x}_{0\mathbf{i}}, h}$ with $\mathbf{i} \in [M]^d$, where $\mathbf{x}_{0\mathbf{i}}$ represents the center of $C_{\mathbf{x}_{0\mathbf{i}}, h}$. Mathematically, we can write

$$C_{\mathbf{x}_{0\mathbf{i}}, h} = \{\mathbf{x} \mid |x_j - x_{0\mathbf{i}, j}| \leq h \text{ for } j \in [d]\}, \quad (3.4)$$

where x_j and $x_{0\mathbf{i}, j}$ are the j^{th} elements of \mathbf{x} and $\mathbf{x}_{0\mathbf{i}}$ respectively. Then the following lemma holds.

Lemma 3.3 (NN Architecture). *Suppose that Assumption 1 holds. There exists a LNN approximate of $g(\mathbf{x})$ as follows:*

$$\|g(\mathbf{x}) - \tilde{s}(\mathbf{x} \mid \tilde{\Lambda})\|_\infty = O(h^p),$$

where $\tilde{\Lambda} = \{\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} \mid \mathbf{i} \in [M]^d\}$, and

$$\tilde{s}(\mathbf{x} \mid \tilde{\Lambda}) = \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i}, h}(\mathbf{x}) \cdot s(\mathbf{x} \mid \mathbf{x}_{0\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}) \quad \text{with} \quad I_{\mathbf{i}, h}(\mathbf{x}) = I(\mathbf{x} \in C_{\mathbf{x}_{0\mathbf{i}}, h}).$$

Remark 3.3.

1. LNN considers the estimation of $\{\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} \mid \mathbf{i} \in [M]^d\}$ simultaneously, where $\tilde{\boldsymbol{\lambda}}_{\mathbf{i}}$ corresponds to $\boldsymbol{\lambda}_{\mathbf{i}}$ up to a rotation matrix \mathbf{D} , and $\boldsymbol{\lambda}_{\mathbf{i}}$ is decided by the Taylor expansion of $g(\mathbf{x})$ at the point $\mathbf{x}_{0\mathbf{i}}$.

2. The use of the indicator function is consistent with the sparsity setting of Schmidt-Hieber (2020), in which the author argues that “*the network sparsity assumes that there are only few non-zero/active network parameters*”. In fact, our study further clarifies the definition of “*the network sparsity*” by providing two categories: (1) defining non-active neurons (such as those in Figure 1) which is realised through the use of indicator function $I_{i,h}(\mathbf{x})$, and (2) pointing out the number of effective parameters. In particular, the second point is important in our view, as one can now define a true set of parameters when using thresholding techniques as in Wang and Lin (2021), and Fan and Gu (2022).
3. The bandwidth h (i.e., the number of cubes) is data driven. This point will become obvious when we move on to Sections 3.2 and 3.3 below.

3.2 On Model (1.1)

In this subsection, we consider Model (1.1) only. As explained previously, LNN targets the entire area that $g(\mathbf{x})$ is defined on, and uses all training data to pre-specify a large number of parameters which do not change with respect to the test set.

In view of Lemma 3.3, we define the following set of LNN candidates:

$$\mathcal{S} = \left\{ \tilde{s}(\mathbf{x} | \Theta) = \sum_{\mathbf{i} \in [M]^d} I_{i,h}(\mathbf{x}) \cdot s(\mathbf{x} | \mathbf{x}_{0\mathbf{i}}, \boldsymbol{\theta}_{\mathbf{i}}) \mid \|\boldsymbol{\theta}_{\mathbf{i}}\| < \infty \right\}, \quad (3.5)$$

in which $\Theta = \{\boldsymbol{\theta}_{\mathbf{i}} \mid \mathbf{i} \in [M]^d\}$ with $\boldsymbol{\theta}_{\mathbf{i}}$'s being $d_q \times 1$ vectors. The objective function is defined as

$$Q(\Theta) = \sum_{t=1}^T [y_t - \tilde{s}(\mathbf{x}_t | \Theta)]^2. \quad (3.6)$$

Accordingly the OLS estimator of $\tilde{\Lambda}$ defined in Lemma 3.3 is obtained by

$$\hat{\Theta} = \underset{\tilde{s} \in \mathcal{S}}{\operatorname{argmin}} Q(\Theta) \quad \text{with} \quad \hat{\Theta} = \{\hat{\boldsymbol{\theta}}_{\mathbf{i}} \mid \mathbf{i} \in [M]^d\} \quad (3.7)$$

and, for $\forall \mathbf{x}_0 \in [-a, a]^d$, the estimator of $g(\mathbf{x}_0)$ is then defined by

$$\hat{g}(\mathbf{x}_0) = \tilde{s}(\mathbf{x}_0 | \hat{\Theta}). \quad (3.8)$$

By design, LNN accounts for the points on the boundary of $[-a, a]^d$ automatically.

Remark 3.4. Our investigation assumes that the number of regressors is correctly specified. When they are over-specified, sparsity naturally arises. It is now even clearer why thresholding methods (such as those in Wang and Lin, 2021; Fan and Gu, 2022) should be adopted, and what

they are really penalizing. In a sense identification restrictions bridge LNN and thresholding approaches in a better fashion. In this work, we do not further explore regression with penalty terms in order not to deviate from our main goal.

Equation (3.7) admits a closed-form estimator for each $\widehat{\boldsymbol{\theta}}_{\mathbf{i}}$. To see this, we write

$$\begin{aligned}\frac{\partial Q(\boldsymbol{\Theta})}{\partial \boldsymbol{\theta}_{\mathbf{i}}} &= -2 \sum_{t=1}^T [y_t - \widetilde{s}(\mathbf{x}_t | \boldsymbol{\Theta})] \cdot I_{\mathbf{i},h}(\mathbf{x}_t) \cdot \frac{\partial s(\mathbf{x}_t | \mathbf{x}_{0\mathbf{i}}, \boldsymbol{\theta}_{\mathbf{i}})}{\partial \boldsymbol{\theta}_{\mathbf{i}}} \\ &= -2 \sum_{t=1}^T [y_t - \widetilde{\mathbf{x}}_{\mathbf{i},t}^\top \boldsymbol{\theta}_{\mathbf{i}}] \cdot \widetilde{\mathbf{x}}_{\mathbf{i},t},\end{aligned}$$

where $\widetilde{\mathbf{x}}_{\mathbf{i},t} = I_{\mathbf{i},h}(\mathbf{x}_t)(\mathbf{I}_{d_q} \otimes \boldsymbol{\gamma}^\top) \boldsymbol{\sigma}(\mathbf{x}_t | \mathbf{x}_{0\mathbf{i}})$, and the second equality follows from the fact that $I_{\mathbf{i},h}(\mathbf{x}_t)I_{\mathbf{j},h}(\mathbf{x}_t) = 0$ for $\mathbf{i} \neq \mathbf{j}$. Thus, for $\forall \mathbf{i}$, the first order condition yields

$$\widehat{\boldsymbol{\theta}}_{\mathbf{i}} = \left(\sum_{t=1}^T \widetilde{\mathbf{x}}_{\mathbf{i},t} \widetilde{\mathbf{x}}_{\mathbf{i},t}^\top \right)^{-1} \sum_{t=1}^T \widetilde{\mathbf{x}}_{\mathbf{i},t} y_t. \quad (3.9)$$

Remark 3.5. We note that by (A.10) and (A.11) of the online supplementary appendix:

$$\sup_{\mathbf{x} \in C_{\mathbf{x}_{0\mathbf{i}},h}} \left\| \mathbf{D}^{-1} \mathbf{m}(\mathbf{x} | \mathbf{x}_{0\mathbf{i}}) - (\mathbf{I}_{d_q} \otimes \boldsymbol{\gamma}^\top) \boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_{0\mathbf{i}}) \right\| = O(h^{q+1}),$$

where \mathbf{D} is a full rank predetermined rotation matrix, and has been explained in Remark 3.3. In connection with the definition of $\widetilde{\mathbf{x}}_{\mathbf{i},t}$, we immediately obtain

$$\left\| \mathbf{D}^{-1} I_{\mathbf{i},h}(\mathbf{x}_t) \cdot \mathbf{m}(\mathbf{x}_t | \mathbf{x}_{0\mathbf{i}}) - \widetilde{\mathbf{x}}_{\mathbf{i},t} \right\| = O_P(h^{q+1}).$$

Thus, $\widetilde{\mathbf{x}}_{\mathbf{i},t}$ is equivalent to $I_{\mathbf{i},h}(\mathbf{x}_t) \cdot \mathbf{m}(\mathbf{x}_t | \mathbf{x}_{0\mathbf{i}})$ up to a predetermined rotation matrix \mathbf{D} . The finding is consistent with Remark 3.2, and ensures the invertibility of $\sum_{t=1}^T \widetilde{\mathbf{x}}_{\mathbf{i},t} \widetilde{\mathbf{x}}_{\mathbf{i},t}^\top$, at least asymptotically.

After carefully studying (3.9) for each \mathbf{i} and repeatedly invoking $I_{\mathbf{i},h}(\mathbf{x}_t)I_{\mathbf{j},h}(\mathbf{x}_t) = 0$ for $\mathbf{i} \neq \mathbf{j}$, the first theorem of this paper is stated as follows.

Theorem 3.1. *Suppose that Assumptions 1 and 2 hold. For $\forall \mathbf{x}_0 \in [-a, a]^d$,*

$$\sqrt{Th^d} \widehat{\boldsymbol{\sigma}}_{\mathbf{x}_0}^{-1} (\widehat{g}(\mathbf{x}_0) - g(\mathbf{x}_0) + O_P(h^p)) \rightarrow_D N(0, 1),$$

where $\widehat{\boldsymbol{\sigma}}_{\mathbf{x}_0}^2 = \sigma_\varepsilon^2 \mathbf{m}(\mathbf{x}_0 | \mathbf{x}_0)^\top \mathbf{H} \boldsymbol{\Sigma}_{\mathbf{x}_0}^{-1} \mathbf{H} \mathbf{m}(\mathbf{x}_0 | \mathbf{x}_0)$, $\boldsymbol{\Sigma}_{\mathbf{x}_0} = f_{\mathbf{x}}(\mathbf{x}_0) \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}$, $\mathbf{H} = \text{diag}\{H_1, \dots, H_{d_q}\}$ with $H_j = \prod_{k=1}^d h^{-n_{j,k}}$, and $\mathbf{n}_j = (n_{j,1}, \dots, n_{j,d})^\top$ includes the corresponding power terms of $m_j(\mathbf{x} | \mathbf{x}_0)$ defined in (3.2).

It is noted that our LNN based estimation method is simple and easy to implement. As pointed out in the introduction, moreover, we are probably among the first being able to rigorously establish the asymptotic normality of the LNN based estimator given by $\widehat{g}(\mathbf{x}_0)$ in Theorem 3.1.

It is also noted out that our simulation studies in Appendix A.2 show that $\widehat{g}(\mathbf{x}_0)$ has better finite-sample properties than those of the conventional local-constant kernel estimation method. The reason to explain the finite-sample advantages is probably that the LNN based estimation method naturally makes the best use of all possible linear combinations of \mathbf{x}_t as explained in Remark 3.2.1 above. By contrast, the conventional local-constant kernel estimation method only evaluates the estimates at the single data points: \mathbf{x}_t .

The same comments apply to Theorem 3.3 below.

Remark 3.6.

1. When establishing the asymptotic distribution, the terms $E[\varepsilon_1\varepsilon_{1+t}]$ for $t \geq 1$ all vanish in the asymptotic covariance matrix due to the partition of (3.4) and the use of the indicator function. Here, the indicator function is a special nonparametric kernel function, so it is not very surprising that the finding is consistent with Chapter 5 of Fan and Yao (2003), for example, in the literature of nonparametric kernel regression with dependent data. As a result, when constructing confidence interval, one can turn to wild bootstrap procedure directly, which also bypasses the rotation matrix \mathbf{D} (see Theorem 3.2 below for details).
2. Ideally, we would like to replace the indicator function by a general kernel function such as $K(\cdot)$ in (3.5). However, by doing so, $s(\mathbf{x} | \mathbf{x}_0, \widetilde{\boldsymbol{\lambda}})$ will not be able to approximate $g(\mathbf{x}_0)$ using Lemma 3.2.
3. We do require $g(\mathbf{x})$ to be defined on a compact set by the design of LNN, but do not impose restriction on the range of $\{\mathbf{x}_t\}$. In Appendix A.3, we further explain how to get a valid result for $\forall \mathbf{x}_0 \in \mathbb{R}^d$ from a nonparametric regression perspective.

To close this subsection, we propose the following bootstrap procedure to establish inference in practice.

1. By (3.8), we calculate $\widehat{\varepsilon}_t = y_t - \widehat{g}(\mathbf{x}_t)$ for $t \in [T]$.
2. Collect i.i.d. draws of $\{\eta_t | t \in [T]\}$ from $N(0, 1)$, and construct the bootstrap version dependent variables as follows: $y_t^* = \widehat{g}(\mathbf{x}_t) + \widehat{\varepsilon}_t\eta_t$. We re-estimate $g(\cdot)$ using $\{(y_t^*, \mathbf{x}_T) | t \in [T]\}$ as in (3.8), and denote the estimate as $\widehat{g}^*(\cdot)$.
3. Repeat Step 2 R times, where R is sufficiently large.

For the bootstrap procedure, the following result holds immediately.

Theorem 3.2. *Let the conditions of Theorem 3.1 hold. Suppose further that $Th^{d+2p} \rightarrow 0$. For $\forall \mathbf{x}_0 \in [-a, a]^d$, we have*

$$\sup_w |\Pr^*(\sqrt{Th^d} \hat{\sigma}_{\mathbf{x}_0}^{-1} [\hat{g}^*(\mathbf{x}_0) - \hat{g}(\mathbf{x}_0)] \leq w) - \Pr(\sqrt{Th^d} \hat{\sigma}_{\mathbf{x}_0}^{-1} [\hat{g}(\mathbf{x}_0) - g(\mathbf{x}_0)] \leq w)| = o_P(1).$$

where \Pr^* is the probability measure induced by the bootstrap procedure.

The condition $Th^{d+2p} \rightarrow 0$ is required to eliminate the bias terms. For the LNN estimators, the bias terms arise from two resources: (1) LNN approximates each polynomial term, and (2) the polynomial terms approximate $g(\cdot)$. As a result, we could not derive the detailed form of the biases but providing its order only. That said, to ensure the bootstrap procedure can provide a valid confidence interval in practice, we require $Th^{d+2p} \rightarrow 0$.

3.3 On Model (1.2)

In this subsection, we investigate Model (1.2), which has been widely adopted for a variety of decision making processes. See Athey (2019) for discussions on different examples.

Direct calculation shows that

$$\begin{aligned} \Pr(y = 1 | \mathbf{x}) &= \Phi_\varepsilon(g(\mathbf{x})), \\ \Pr(y = 0 | \mathbf{x}) &= 1 - \Phi_\varepsilon(g(\mathbf{x})), \end{aligned}$$

which yields $E[y | \mathbf{x}] = \Phi_\varepsilon(g(\mathbf{x}))$. Thus, provided a set of time series observations given in (1.3), the likelihood function is specified as follows:

$$L(g) = \prod_{t=1}^T [1 - \Phi_\varepsilon(g(\mathbf{x}_t))]^{1-y_t} \cdot \Phi_\varepsilon(g(\mathbf{x}_t))^{y_t}.$$

Accordingly, the log-likelihood function is defined below:

$$\begin{aligned} \log L(g) &= \sum_{t=1}^T l_t(g(\mathbf{x}_t)) \\ &= \sum_{t=1}^T \{(1 - y_t) \cdot \log[1 - \Phi_\varepsilon(g(\mathbf{x}_t))] + y_t \cdot \log \Phi_\varepsilon(g(\mathbf{x}_t))\}, \end{aligned} \quad (3.10)$$

where the definition of $l_t(\cdot)$ is obvious.

Our interest is still inferring $g(\cdot)$. We construct the neural network as in (3.5), and consider the following objective function:

$$\log L(\tilde{s}(\cdot | \Theta)) = \sum_{t=1}^T \log l_t(\tilde{s}(\mathbf{x}_t | \Theta))$$

$$= \sum_{t=1}^T \{(1 - y_t) \cdot \log[1 - \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta))] + y_t \cdot \log \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta))\},$$

which yields the following maximum likelihood estimator:

$$\hat{\Theta} = \operatorname{argmax}_{\tilde{s} \in \mathcal{S}} \log L(\tilde{s}(\cdot | \Theta)). \quad (3.11)$$

Note that our LNN method involves a general unknown function form. As a result, the classical results of likelihood estimation, such as those in Newey and McFadden (1994), no longer hold. Therefore, before establishing an asymptotic distribution using $\hat{\Theta}$, we state a lemma to show the overall feasibility of the LNN architecture when modelling binary outcomes.

Lemma 3.4. *Under Assumptions 1 and 2,*

$$\frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(g(\mathbf{x}_t)) - \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \hat{\Theta}))]^2 = o_P(1).$$

Lemma 3.4 provides the consistency, and also bridges the likelihood estimation and the nonlinear least squares approach to some extent. To be precise, Lemma 3.4 does not provide any specific consistency for $\forall \hat{\theta}_i$. Instead, it evaluates the overall performance of LNN. More importantly, it says when modelling a binary outcome, the likelihood estimation using the LNN architecture is approximately equivalent to implementing a nonlinear least squares method provided the distribution of ε_t is correctly specified. In addition, Lemma 3.4 further infers that

$$\frac{1}{T} \sum_{t=1}^T [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \hat{\Theta})]^2 = o_P(1),$$

which is nice, as many remarks made in Section 3.2 can be directly applied. Last but not least, Lemma 3.4 facilitates numerical implementation in practice, which will be further discussed in Section 4.2.

Below, we establish the following asymptotic distribution for Model (1.2).

Theorem 3.3. *Suppose that Assumptions 1 and 2 hold. For $\forall \mathbf{x}_0 \in [-a, a]^d$,*

$$\sqrt{Th^d \tilde{\sigma}_{\mathbf{x}_0}^{-1}} (\hat{g}(\mathbf{x}_0) - g(\mathbf{x}_0) + O_P(h^p)) \rightarrow_D N(0, 1),$$

where $\hat{g}(\mathbf{x})$ is defined in the same form as (3.8), $\tilde{\sigma}_{\mathbf{x}_0}^2 = \mathbf{m}(\mathbf{x}_0 | \mathbf{x}_0)^\top \mathbf{H} \tilde{\Sigma}_{\mathbf{x}_0}^{-1} \mathbf{H} \mathbf{m}(\mathbf{x}_0 | \mathbf{x}_0)$, and

$$\tilde{\Sigma}_{\mathbf{x}_0} = \frac{f_{\mathbf{x}}(\mathbf{x}_0) \phi_\varepsilon(g(\mathbf{x}_0))^2}{[1 - \Phi_\varepsilon(g(\mathbf{x}_0))] \Phi_\varepsilon(g(\mathbf{x}_0))} \int_{[-1, 1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}.$$

Remark 3.6 can be applied to Theorem 3.3 with some obvious modifications.

In light of Patrick and Andres (2012), we then propose a score based wild bootstrap approach for inferential purposes as follows.

1. For each bootstrap replication, we collect i.i.d. draws of $\{\eta_t \mid t \in [T]\}$ from $N(0, 1)$, and calculate

$$\widehat{\boldsymbol{\theta}}_{\mathbf{i}}^* = \widehat{\boldsymbol{\theta}}_{\mathbf{i}} + \left(\sum_{t=1}^T \frac{\partial^2 \log l_t(\tilde{s}(\mathbf{x}_t \mid \widehat{\boldsymbol{\Theta}}))}{\partial \boldsymbol{\theta}_{\mathbf{i}} \partial \boldsymbol{\theta}_{\mathbf{i}}^\top} \right)^{-1} \sum_{t=1}^T \frac{\partial \log l_t(\tilde{s}(\mathbf{x}_t \mid \widehat{\boldsymbol{\Theta}}))}{\partial \boldsymbol{\theta}_{\mathbf{i}}} \eta_t, \quad (3.12)$$

where $l_t(\cdot)$ is defined in (3.10).

2. Repeat Step 1 R times, where R is sufficiently large.

It is worth pointing out that the above procedure is computationally efficient in the sense that the right hand side of (3.12) enjoys a closed-form expression, which is given in (A.28) and (A.29) specifically for the sake of space. Practically, the bootstrap procedure may require much less time compared with the estimation of (3.11) itself. We will further comment on the computational issues associated with Model (1.2) in Section 4.2.

The following theorem holds for the above bootstrap procedure.

Theorem 3.4. *Let the conditions of Theorem 3.3 hold. Suppose further that $Th^{d+2p} \rightarrow 0$. For $\forall \mathbf{x}_0 \in [-a, a]^d$, we have*

$$\sup_w |\Pr^*(\sqrt{Th^d} \tilde{\sigma}_{\mathbf{x}_0}^{-1} [\widehat{g}^*(\mathbf{x}_0) - \widehat{g}(\mathbf{x}_0)] \leq w) - \Pr(\sqrt{Th^d} \tilde{\sigma}_{\mathbf{x}_0}^{-1} [\widehat{g}(\mathbf{x}_0) - g(\mathbf{x}_0)] \leq w)| = o_P(1),$$

where \Pr^* is the probability measure induced by the bootstrap procedure, and $\widehat{g}^*(\cdot)$ is yielded by the bootstrap draws in an obvious manner.

The comments and discussion below Theorem 3.2 apply to Theorem 3.4 with some minor modifications, so we do not repeat them here.

3.4 Further Discussion

Up to this point, we would like to point out that those questions raised in Section 1 have all been answered, so Figure 1 can be understood better. We summarize some key points which may have been discussed previously here and there, and further discuss some issues left behind.

Neurons — Having established the results in Sections 3.2 and 3.3, it is now clear that for LNN, the total number of neurons is

$$M^d \cdot d_q(q+1) = \left(\frac{2a}{h} \right)^d \cdot d_q(q+1),$$

of which only $d_q(q + 1)$ neurons are activated when loading test data. Among the activated neurons, the number of effective parameters is only d_q , while the rest of the parameters are predetermined.

Multiple Hidden Layers — Lemma 3.1 yields a recursive relationship, as one can repeatedly invoke Lemma 3.1 to replace $(x - x_0)$ inside the activation. It will then yield a LNN architecture with multiple hidden layers. Although having multiple hidden layers is achievable, at this stage it is not clear to us why we should do so. As discussed in Remark 3.1, this step is completely independent of data, so we do not see any benefit of doing so unless $g(\cdot)$ in both models has certain specific structure. Under some extra structure on $g(\cdot)$, however, the necessity of developing a LNN architecture with multiple hidden layers deserves extra attention in future research.

Dependence — As shown for both models, LNN automatically eliminates the correlation of observations from different time periods when establishing the asymptotic distribution. This finding is important, as one can easily use different types of wild bootstrap methods to obtain valid inference.

ReLU Function — The current study focuses on Sigmoidal activation function. We conjecture for ReLU function a corresponding LNN approach can be constructed similarly, which, however, requires some careful consideration.

Thresholding — Our investigation assumes that the number of regressors is correctly specified. When they are over-specified, sparsity naturally arises. In this case, it is clear that why a thresholding method (such as those in Wang and Lin, 2021; Fan and Gu, 2022) should be adopted, and what it is really penalizing. Utilizing identification conditions allows one to define the true set of parameters, so dimension reduction techniques such as LASSO (Tibshirani, 1996) and bridge estimation (Huang et al., 2008) techniques can be adopted straight away.

Trending — In this paper, we assume that the regressors $\{\mathbf{x}_t \mid t \in [T]\}$ are strictly stationary and mixing. In fact, they can have other more complex structures, such as linear processes, locally stationarity, heterogeneity, deterministic trends, etc. As a result, many climate models (such as those in Mudelsee, 2019) may be better captured. For such cases, one may need to revise the assumptions and proofs accordingly depending on detailed research questions.

4 Simulation

In this section we conduct simulations to examine the theoretical findings. We consider Models (1.1) and (1.2) in Sections 4.1 and 4.2 respectively.

Before proceeding further, we would like to point out the following simulation design focuses on coverage rates for the entire test set. Also, no statistical software package is required, as the algorithm becomes obvious after presenting Section 3. We now provide some details of the numerical implementation. As shown in Section 3, many parameters are involved in the LNN architecture. It would be extremely difficult to systematically check every single one in one paper, so we have to be selective. That said, the following quantities are pre-fixed without loss of generality.

- Throughout, we use the Sigmoidal squasher, $\sigma(w) = 1/(1 + \exp(-w))$, as the activation function. The derivatives of Sigmoidal squasher are provided in Appendix A.1 of the online supplementary file for the sake of space.
- $\boldsymbol{\pi}_j$'s are generated in exactly the same way as mentioned in Remark 3.2.
- Let $R = 200$ for the bootstrap procedure.
- Let $g(\mathbf{x}) = 1 + \sin(\mathbf{x}^\top \mathbf{1}_d/d)$.

Surely, we can explore different settings, but these are less important based on the development of Section 3. Also, we are constrained by computing power.

4.1 Model (1.1)

Consider the following regression model:

$$y_t = g(\mathbf{x}_t) + \varepsilon_t, \quad (4.1)$$

where $\varepsilon_t = 0.5\varepsilon_{t-1} + N(0, 0.75)$, and the j^{th} element of \mathbf{x}_t is generated as $x_{t,j} \sim U(-a, a)$. The bandwidth h is set as $h = a/M$, where M is the integer closest to a/h_1 with $h_1 = 2.5 \cdot T^{-1/(d+2p-0.5)}$. Here, -0.5 is to ensure $\sqrt{T}h^{p+d/2} \rightarrow 0$ holds. In fact, h is very close to h_1 , and the current setup is simply to guarantee M is a large positive integer. When designing the simulations, our impression is that the results are not sensitive to the choices of $g(\mathbf{x})$ and the bandwidth. Therefore, in what follows, we only vary the values of

$$(q, d, u_\sigma). \quad (4.2)$$

To measure the finite sample performance, we select L^d test points from $[-a, a]^d$ as follows:

$$\mathbf{x}_{L,\mathbf{j}} = \left(-a + \frac{2a}{L-1}(j_1 - 1), \dots, -a + \frac{2a}{L-1}(j_d - 1) \right)^\top, \quad (4.3)$$

where $\mathbf{j} \in [L]^d$. With each dataset, we first estimate all $g(\mathbf{x}_{L,\mathbf{j}})$ using the approach of Section 3.2, and then construct the corresponding 95% confidence interval using the bootstrap procedure documented in Theorem 3.2 for each point. After n replications, we report

$$\text{RMSE}_g = \left\{ \frac{1}{nL^d} \sum_{i=1}^n \sum_{\mathbf{j} \in [L]^d} [\hat{g}_i(\mathbf{x}_{L,\mathbf{j}}) - g(\mathbf{x}_{L,\mathbf{j}})]^2 \right\}^{1/2}$$

and

$$\text{CR}_g = \frac{1}{nL^d} \sum_{i=1}^n \sum_{\mathbf{j} \in [L]^d} I(\hat{g}_i(\mathbf{x}_{L,\mathbf{j}}) - g(\mathbf{x}_{L,\mathbf{j}}) \in \text{CI}_{i,\mathbf{j}}),$$

where $\hat{g}_i(\cdot)$ stands for the estimate of $g(\cdot)$ at the i^{th} replication, and $\text{CI}_{i,\mathbf{j}}$ is the 95% confidence interval of $\hat{g}_i^*(\mathbf{x}_{L,\mathbf{j}}) - \hat{g}_i(\mathbf{x}_{L,\mathbf{j}})$ based on the bootstrap draws from the i^{th} replication. We set $T \in \{800, 1600, 2400\}$ to examine the convergence of NN approach, and let² $u_\sigma \in \{-0.5, 0.5\}$ and $d \in \{2, 3\}$. We further set $a = 3$, $L = 20$ and $n = 200$ without loss of generality.

We start presenting the results below. First, we draw some plots for the case with $d = 2$. In both Figures 2 and 3, the first sub-plot is always the true $g(\mathbf{x})$. For the rest of sub-plots, each has three layers. The middle one is the average of estimates over n replications. The top and bottom layers are the averages of the bootstrap draws corresponding to the 97.5% and 2.5% quantiles respectively over n replications. A few facts emerge. Overall, the LNN approach can recover the unknown function reasonably well. When $q = 4$, both figures have smoother plots, which should be expected provided q is considered as fixed. Also, both figures are very similar, so the results are not sensitive to the choice of u_σ as explained in Remark 3.1.

More detailed numbers are summarized Table 1. As expected, when T goes up, RMSE_g converges to 0 and CR_g converges to 0.95. In particular, for CR_g , when the sample size is not large enough, we tend to get a narrower CI rather than wider CI. Also, we note that when d increases, RMSE_g increases rather significantly, which is understood as the curse of dimensionality. The curse is also reflected in CR_g , but the results are still acceptable in our view. Recall that Figures 2 and 3 suggest that when $q = 4$ the plots are much smoother, but Table 1 actually says larger q is corresponding to larger RMSE_g . However, the change in q does not alter CR_g much, which shows certain robustness of the bootstrap procedure suggested in Section 3.2. Finally, we note that the time consuming part of this simulation design comes from generating inference for a large number of points defined in (4.3) using the proposed bootstrap procedure.

²We provide some extra simulation results with larger d value in the online supplementary appendices of the paper, and comment on some issues related to computational overhead.

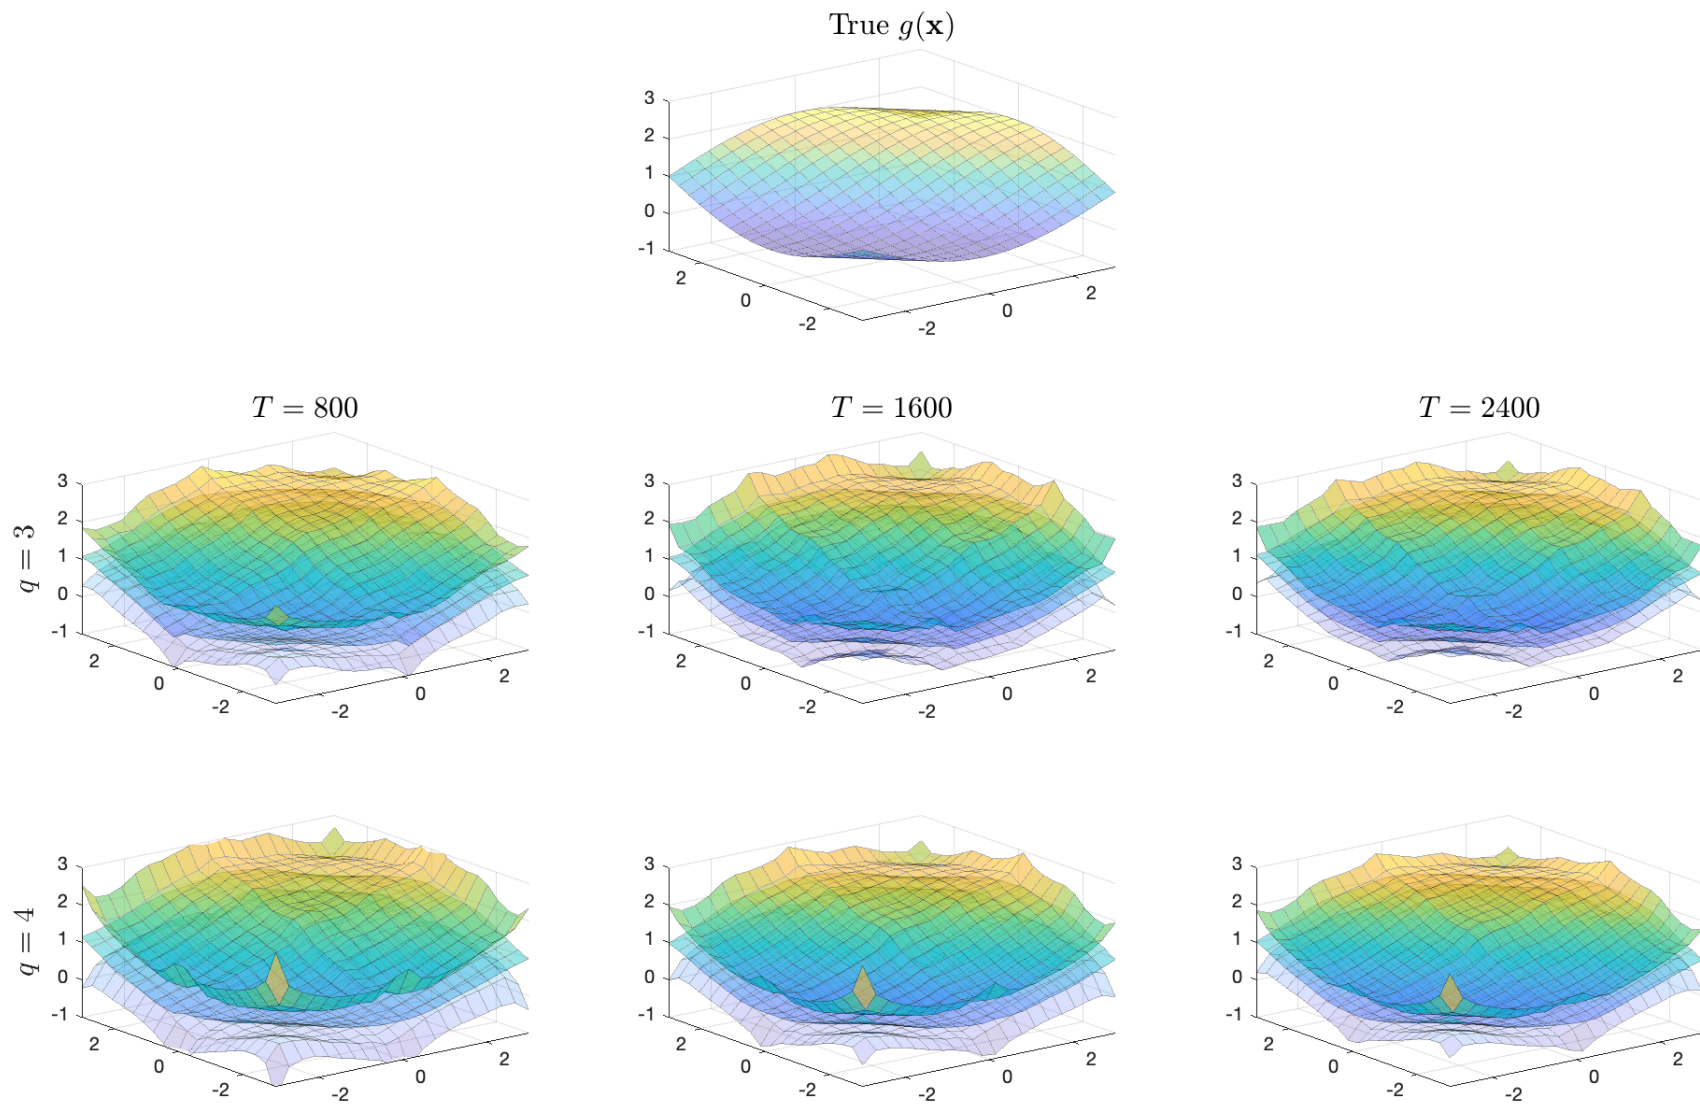


Figure 2: Simulation Results of Example 1.1 ($u_\sigma = -0.5, d = 2$)

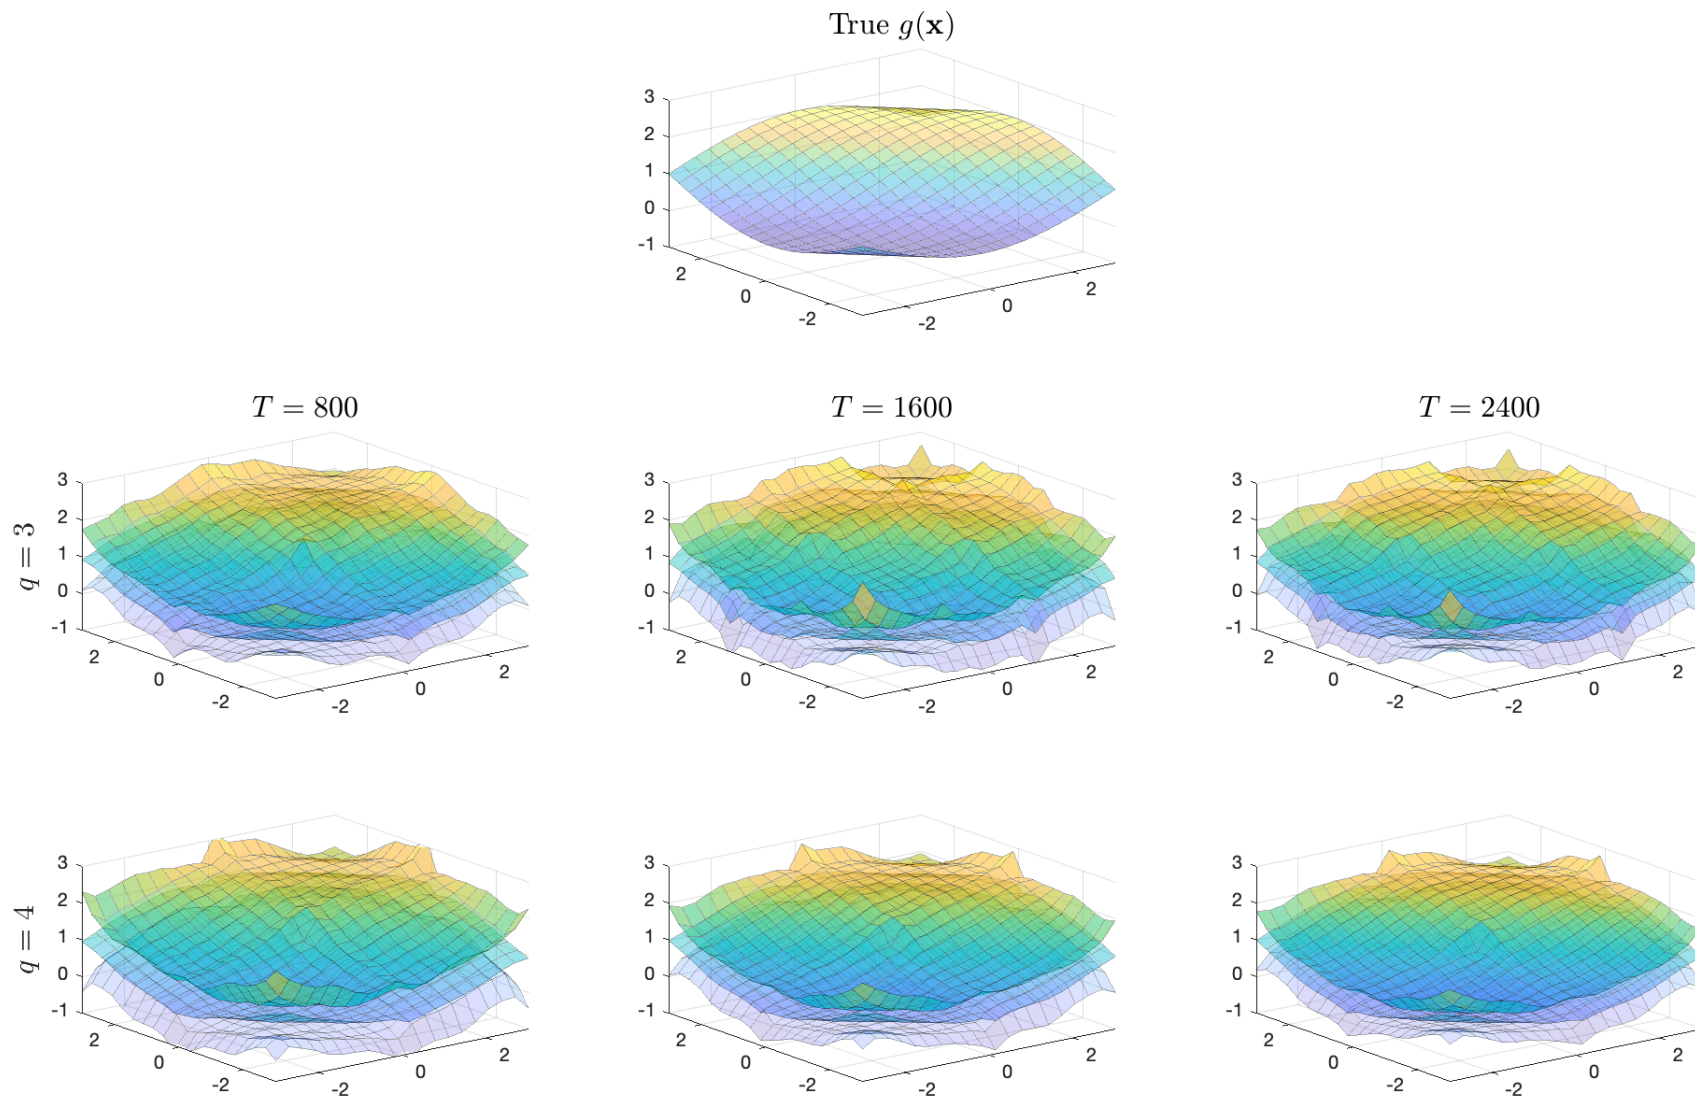


Figure 3: Simulation Results of Example 1.1 ($u_\sigma = 0.5, d = 2$)

Table 1: Simulation Results of Example 1.1

				RMSE _g		CR _g	
		$T \setminus d$	2	3	2	3	
$u_\sigma = -0.5$	$q = 3$	800	0.282	0.637	0.894	0.871	
		1600	0.261	0.543	0.925	0.884	
		2400	0.211	0.356	0.929	0.916	
	$q = 4$	800	0.324	0.932	0.917	0.837	
		1600	0.222	0.529	0.927	0.896	
		2400	0.180	0.408	0.934	0.916	
$u_\sigma = 0.5$	$q = 3$	800	0.291	0.608	0.912	0.873	
		1600	0.261	0.388	0.917	0.907	
		2400	0.236	0.293	0.918	0.910	
	$q = 4$	800	0.327	0.126	0.915	0.837	
		1600	0.226	0.579	0.922	0.891	
		2400	0.184	0.433	0.927	0.911	

4.2 Model (1.2)

Next, we consider the following data generating process:

$$y_t = \begin{cases} 1 & g(\mathbf{x}_t) - \varepsilon_t \geq 0 \\ 0 & \text{otherwise} \end{cases},$$

in which $\{\mathbf{x}_t\}$, $\{\varepsilon_t\}$ and $g(\cdot)$ are exactly the same as those in Section 4.1.

We first note a computational issue. We rely on “fminunc” function of Matlab to find the solution of

$$\operatorname{argmin}_{\tilde{s} \in \mathcal{S}} [-\log L(\tilde{s}(\cdot) | \Theta)],$$

which is the same as that in (3.11). In order to invoke the minimization process in any statistical software (including R, Matlab, etc.), one needs to provide initial values to the parameters under estimation. As a consequence, the numbers reported below are affected by the initial values more or less. Although it is not our intention to tackle this complicated computational issue in this paper, Lemma 3.4 does become useful in this case. Recall that Lemma 3.4 bridges the log likelihood estimation and the nonlinear least squares estimation. Therefore, we first conduct an OLS estimation using the approach of Section 4.1 as the initial value of Θ for each generated $\{(y_t, \mathbf{x}_t) | t \in [T]\}$. We then invoke log likelihood estimation as our final estimate of

Θ for each dataset. Even in this case, the computation is rather slow, and the computational time increases dramatically when the number of parameters goes up.

That said, due to the limitation of computing power, we consider $d = 2$ only in this section. The rest settings are the same as those in Section 4.1. On top of the criteria RMSE_g and CR_g introduced above, we further introduce

$$\text{RMSE}_g^* = \left\{ \frac{1}{\#\mathcal{Q}_{L,\mathbf{j}} \cdot L^d} \sum_{\mathbf{j} \in [L]^d} \sum_{\widehat{g}_i(\mathbf{x}_{L,\mathbf{j}}) \in \mathcal{Q}_{L,\mathbf{j}}} [\widehat{g}_i(\mathbf{x}_{L,\mathbf{j}}) - g(\mathbf{x}_{L,\mathbf{j}})]^2 \right\}^{1/2},$$

where $\mathcal{Q}_{L,\mathbf{j}}$ defines the values between 1st and 3rd quartiles of $\{\widehat{g}_i(\mathbf{x}_{L,\mathbf{j}}) \mid i \in [n]\}$. By doing so, we can remove some impacts of the initial value issues associated with “fminunc”. A similar treatment can also be found in the simulation study of Bauer and Kohler (2019), in which some quantiles of the simulation results are reported in order to eliminate the impacts of outliers.

Again, we draw a few plots as in Figures 2 and 3. In this case, to minimize the impacts of initial value issues caused by the program, we take the median based on the simulation results when plotting each layer. As shown in Figures 4 and 5 below, the results are clearly not as good as those for Model 1.1 due to lack of closed-form solutions. Still, we are able to see the consistency as the sample size increases. Both figures show that the patterns are very much similar to those in Figures 2 and 3, so we omit the discussion here.

We further summarize the detailed numbers in Table 2. A few facts should be mentioned. First, the coverage rates are reasonably well. The values of RMSE_g^* are comparable to those reported in Table 1, in which the numbers are generated from closed-form estimates. Third, the results are not changing much with respect to the value of u_σ , which again confirms our argument in Remark 3.1.

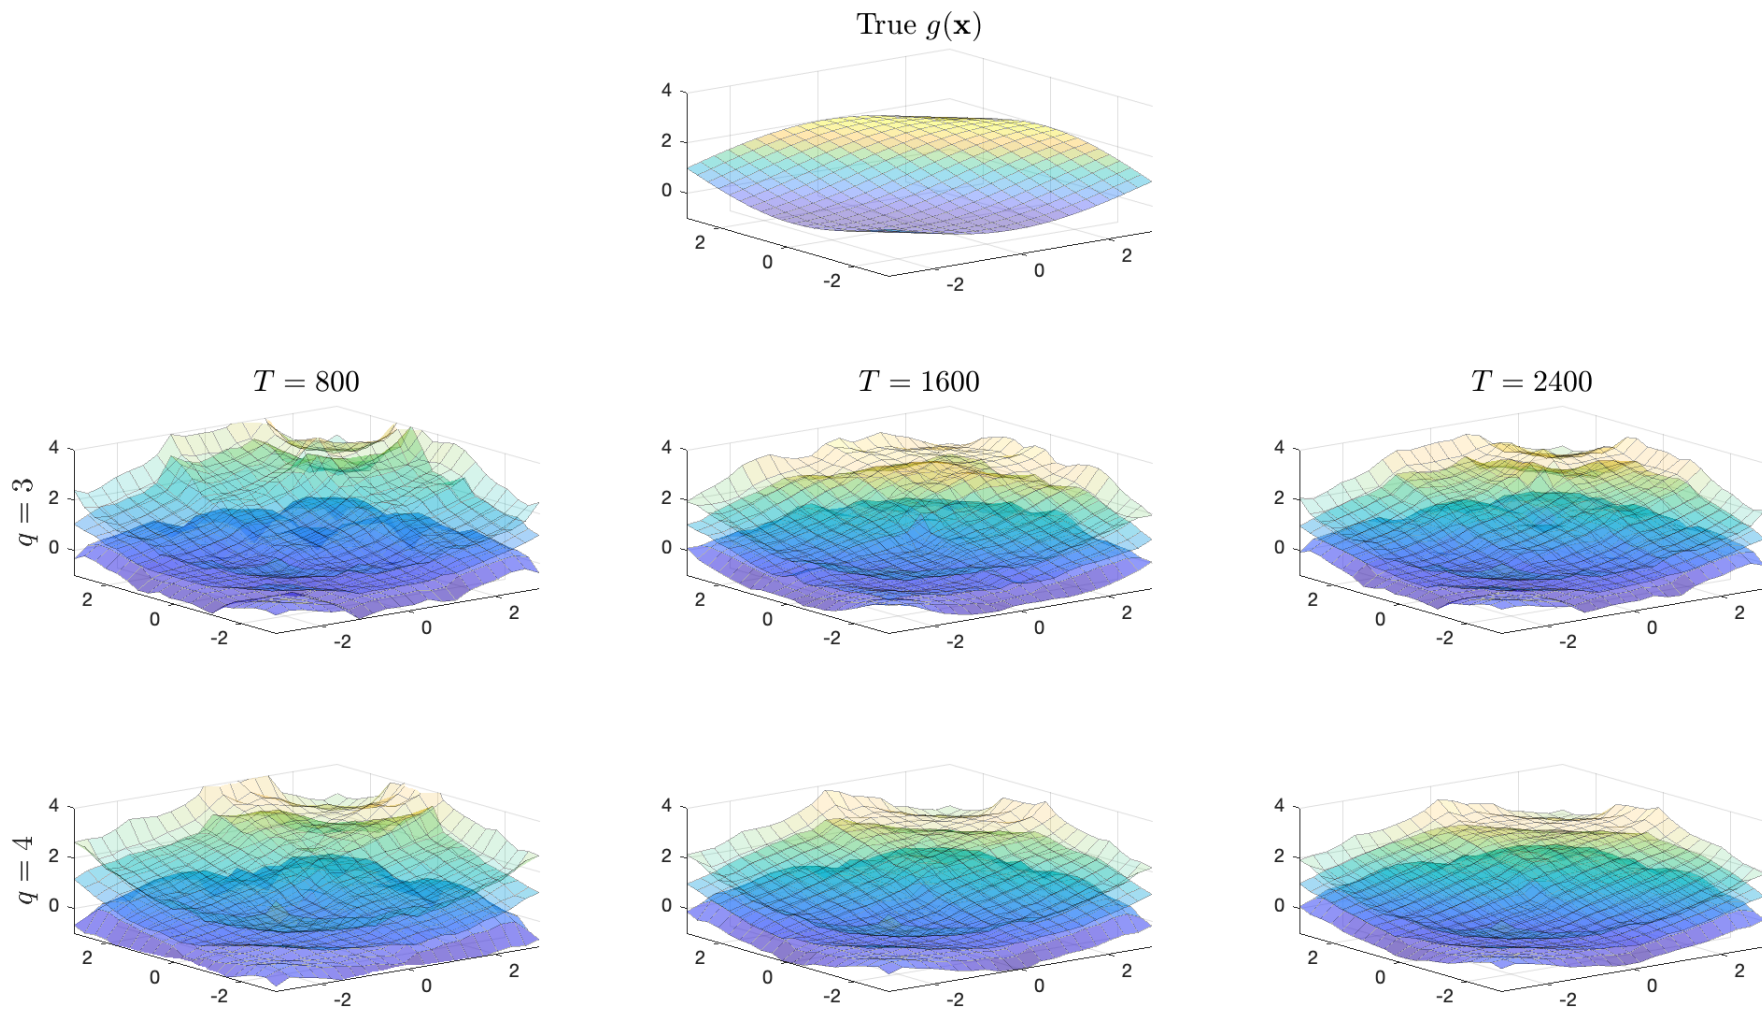


Figure 4: Simulation Results of Example 1.2 ($u_\sigma = -0.5, d = 2$)




Figure 5: Simulation Results of Example 1.2 ($u_\sigma = 0.5, d = 2$)

Table 2: Simulation Results of Example 1.2

		T	RMSE $_g$	RMSE $_g^*$	CR $_g$
$u_\sigma = -0.5$	$q = 3$	800	1.212	0.790	0.904
		1600	1.042	0.537	0.912
		2400	0.967	0.428	0.917
	$q = 4$	800	1.378	0.813	0.915
		1600	0.797	0.405	0.923
		2400	0.525	0.282	0.934
$u_\sigma = 0.5$	$q = 3$	800	1.106	0.600	0.898
		1600	1.023	0.577	0.905
		2400	0.795	0.449	0.907
	$q = 4$	800	1.110	0.616	0.906
		1600	0.687	0.384	0.919
		2400	0.480	0.278	0.930

5 Empirical Study

In this section, we consider two empirical examples to illustrate the two models that have been investigated in Section 3.

5.1 On Climate Data

First, we would like to recall that as mentioned in Section 2, we assume that $\{\mathbf{x}_t \mid t \in [T]\}$ are strictly stationary for the purpose of simplicity. In practice, it is not surprising that data may present certain time trend or seasonality as we shall show below. Such features do not alter our theoretical results fundamentally. For example, we can suppose that

$$\mathbf{x}_t = \mathbf{u}(\tau_t) + \mathbf{s}_t + \mathbf{v}_t, \quad (5.1)$$

where $\tau_t = \frac{t}{T}$, $\mathbf{u}(\cdot)$ is a vector of deterministic trending functions, \mathbf{s}_t captures some seasonal effects, and \mathbf{v}_t mimics randomness. With some obvious modifications on Assumption 2, all the results established previously still hold.

The data to be investigated are collected from Weather Underground API, and has also been extensively studied by researchers from a very diversified background on Kaggle³. This dataset includes four time series (temperature, humidity, wind speed, atmospheric pressure)

³See <https://www.kaggle.com/datasets/sumanthvrao/daily-climate-time-series-data?resource=download> for details.

from 1st January 2013 to 24th April 2017 for the city of Delhi, India. As a routine exercise, we normalize each time series to get sample mean 0 and sample standard deviation 1 respectively. Typically, researchers use the time period from 1st January 2017 to 24th April 2017 as the test data, and use the rest time period as the training data. We do the same, so we have $T = 1461$ observations to train our LNN, and the size of test data is $T^* = 114$.

We plot the normalized time series in Figure 6 below. In each sub-plot, the observations on the left hand of the vertical line are the training data, while the observations on the right hand side are the test data. Temperature and pressure move in opposite direction, while humidity and wind speed look like white noises more or less. Although wind speed has some spikes from time to time, it is rather stationary overall. When training our LNN, we consider Model 1.1 and fit temperature at period t (i.e., y_t) on humidity, wind speed and atmospheric pressure at period $t - 1$ (i.e., \mathbf{x}_t).

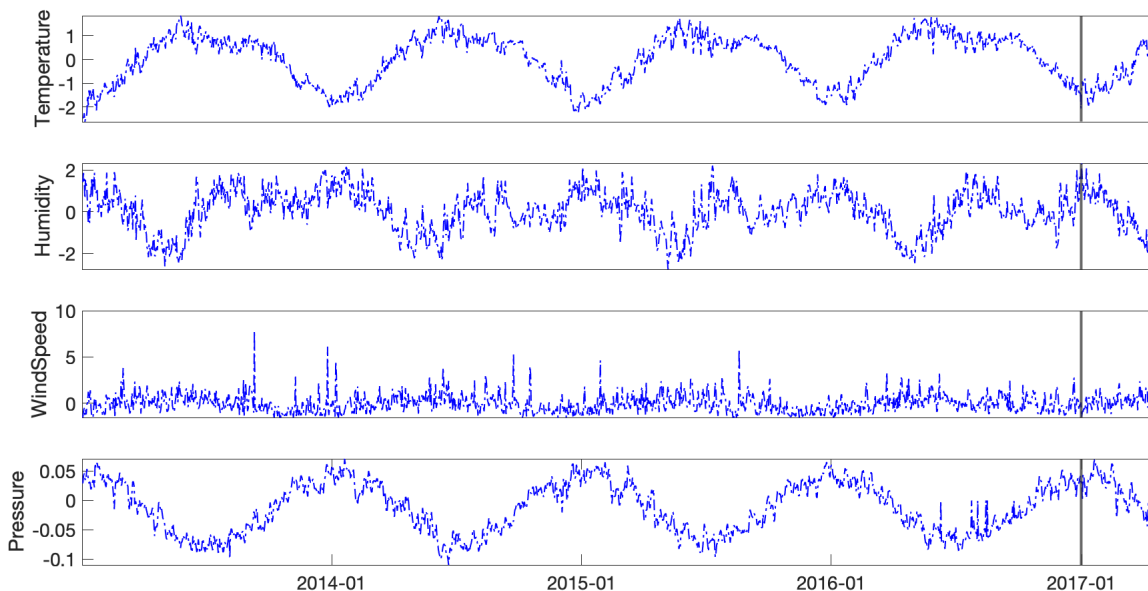


Figure 6: Climate Data of Delhi

In view of the data range, we specifically consider the following settings when running the regression:

$$a = 2, \quad q \in \{3, 4, 5, 6\}, \quad u_\sigma \in \{-0.5, 0.5\} \quad (5.2)$$

The rest settings are the same as those in simulation studies. First, we report RMSE using test set. Specifically, after training our LNN, we calculate

$$\text{RMSE} = \left\{ \frac{1}{T^*} \sum_{t=1}^{T^*} [y_t - \hat{g}(\mathbf{x}_t)]^2 \right\}^{1/2}, \quad (5.3)$$

where (y_t, \mathbf{x}_t) are from the test set, and $\hat{g}(\cdot)$ is obtained using the training set.

We summarize RMSE in Table 3 below, in which the results are very close. Regardless the value of u_σ , when we increase the value of q , RMSE always shows a U shape, which can be used as a criterion to pick the “optimal” q in practice. In what follows, we focus on the cases $(u_\sigma, q) = (0.5, 5)$ and $(u_\sigma, q) = (-0.5, 4)$, as both combinations are corresponding to the bottom points of each U-curve.

In both Figures 7 and 8, the black dotted line stands for the true data, the red dashdotted line stands for the estimated values, and the blue solid lines present the 95% CI based on the bootstrap draws. Figure 7 shows that the LNN approach can capture the trend of the entire test period reasonably well over all, and Figure 8 provides a zoomed-in version by focusing on the test period only. It is clear that both subplots are almost identical, although the values of (u_σ, q) are different. Therefore, we conclude that the numerical results show certain robustness of NN architecture.

Table 3: RMSE for the Test Set

$u_\sigma \setminus q$	3	4	5	6
0	0.5447	0.5227	0.5287	0.5354
0.5	0.5719	0.5267	0.5238	0.5368




Figure 7: Plot of the Fitted Curve on the Entire Period

5.2 On Clean Energy Index

As pointed by Christoffersen and Diebold (2006) and many follow-up studies, the sign of stock market returns may be predictable even if the returns themselves are not predictable. In this section, we specifically use Model 1.2 to investigate the clean energy index (ECO) assembled by




Figure 8: Plot of the Fitted Curve of the Test Period

WiderShares, LLC, which is a modified equal weighted index comprised of companies that are publicly traded in the United States and that are engaged in the business of the advancement of cleaner energy and conservation. A very detailed data description about ECO can be found at <https://wildershares.com/about.php>, so we omit them here for simplicity. In connection with the recent trend of investigating climate change, the importance of understanding clean energy sector is more important than ever. Not to mention that for finance discipline, studying energy sector is always a topical interest (e.g., Jin and Jorion, 1996; El-Sharif et al., 2005; and many follow-up studies).

In what follows, we focus on the next model

$$y_t = \begin{cases} 1 & g(\mathbf{x}_t) - \varepsilon_t \geq 0 \\ 0 & \text{otherwise} \end{cases},$$

where $y_t = 1$ if ECO index return is positive at time t , otherwise $y_t = 0$; and \mathbf{x}_t includes oil price change rate at time $t - 1$, gas price change rate at time $t - 1$, and CBOE volatility index (VIX) change rate at time $t - 1$. The choices of regressors are based on Section III of Jin and Jorion (1996) and Section 7 of Christoffersen and Diebold (2006). One definitely can consider more choices and different forms of regressors, which may lead to another research paper with comprehensive empirical study. Again, we normalize the observations of each regressor to ensure mean 0 and standard deviation 1. Finally, we present the data in Figure 9.

The entire time period we consider covers 4th January 2010 to 31st December 2019, in which we use the period from 1st July 2019 to 31st December 2019 as the test period. As a result, we have 2369 observations as training data (i.e., $T = 2369$), and 146 observations as test data (i.e., $T^* = 146$).




Figure 9: Clean Energy Index Data

After training our LNN, we calculate the mean absolute error (MAE) as follows:

$$\text{MAE} = \frac{1}{T^*} \sum_{t=1}^{T^*} |y_t - \hat{y}_t|,$$

where $\hat{y}_t = 1$ if $\hat{g}(\mathbf{x}_t) \geq 0.5$, and $\hat{y}_t = 0$ otherwise. Here, we arbitrarily use 0.5 as the threshold. One may further explore other options in practice. We consider different choices of u_σ and q , and summarize results in Table 4 below. Regardless the value of u_σ , we obtain the minimum MAE with $q = 4$. In the best case scenario, the successful rate is around 59% which might be further improved by accounting for volatility as well as other possible explanatory variables through a comprehensive empirical investigation.

Table 4: MAE for the Test Set

$u_\sigma \setminus q$	3	4	5
-0.5	0.5068	0.4384	0.5137
0.5	0.5205	0.4110	0.4589

In what follows, we focus on the results by setting $q = 4$. In Figure 10, using test data we respectively plot the estimated $\hat{g}(\cdot)$, and the corresponding probabilities (i.e., $\Phi_\varepsilon(\hat{g}(\cdot))$). In each sub-figure, the dash-dotted line is corresponding to the choice of $(u_\sigma = -0.5, q = 4)$, while the

dotted line is corresponding to the choice of $(u_\sigma = 0.5, q = 4)$. The results are almost identical when u_σ varies, so it again justifies our argument on the choice of u_σ and the robustness of the approach investigated above.




Figure 10: Plot of $\hat{g}(\mathbf{x})$ and associated Probabilities of the Test Period

Finally, we point out that Christoffersen and Diebold (2006) write “*As volatility moves, so too does the probability of a positive return: the higher the volatility, the lower the probability of a positive return*”. With good sign prediction, we can then remove the “*bad volatility*” and keep the “*good volatility*” (e.g., Patton and Sheppard, 2015). This is important, as traditionally a primary goal of portfolio analysis is to minimize the volatility (e.g., Engle et al., 2019), so many good volatilities are wiped out by classic approaches. Model (1.2) associated with the LNN approach naturally marries the above studies and improves the prediction of positive returns, so a better portfolio strategy can be implemented in practice. We will no longer proceed further along this line, and would like to leave this in another research paper by systematically comparing different approaches adopted when constructing portfolio.

6 Conclusion

NN has gained considerable attentions over the past a few decades. Yet, questions as those stated in Section 1 have not been answered well in the literature. In this paper, we bring in identification restrictions to the LNN framework from a nonparametric regression perspective, and consider the LNN based estimation with dependent data for both cases with quantitative/qualitative outcomes. We then establish the LNN based estimation theory under a set of minor conditions. The asymptotic distributions are derived accordingly, and we show that LNN automatically eliminates the dependence of data when calculating the asymptotic variances.

The finding is important, as one can easily use different types of wild bootstrap methods to obtain valid inference for practical implementation. In particular, for quantitative outcomes, the LNN approach yields closed-form expressions for the estimates of some key parameters of interest. Last but not least, we examine our theoretical findings through extensive numerical studies.

Several major comments have been made here and there in Section 3, and some future research directions have been acknowledged along the way. Finally, we hope the current article will shed light on how to produce transparent algorithms to ensure that our research findings are relevant and useful for practical implementations and applications.

7 Acknowledgements

Gao, Peng and Yang acknowledge financial support from the Australian Research Council Discovery Grants Program under Grant Numbers: DP200102769, DP210100476 and DP230102250, respectively.

References

- Abramovitz, M. and Stegun, I. A. (1972), *Handbook of mathematical functions*, Dover Publications, New York, U.S.
- Athey, S. (2019), The impact of machine learning on economics, in J. G. Ajay Agrawal and A. Goldfarb, eds, ‘The Economics of Artificial Intelligence: An Agenda’, pp. 507–547.
- Bartlett, P. L., Montanari, A. and Rakhlin, A. (2021), ‘Deep learning: A statistical viewpoint’, *Acta Numerica* **30**, 87201.
- Bauer, B. and Kohler, M. (2019), ‘On deep learning as a remedy for the curse of dimensionality in nonparametric regression’, *The Annals of Statistics* **47**(4), 2261–2285.
- Chen, X. (2007), ‘Large Sample Sieve Estimation of Semi-Nonparametric Models, Chapter 76 edited by James J. Heckman and Edward E. Leamer’, *Handbook of Econometrics* **6B**, 5549–5632.
- Chen, X., Liu, Y., Ma, S. and Zhang, Z. (2022), Casual inference of general treatment effects using neural networks with a diverging number of confounders. Available at <https://arxiv.org/abs/2009.07055v5>.
- Chen, X., Racine, J. and Swanson, N. (2001), ‘Semiparametric ARX neural network models with an application to forecasting inflation’, *IEEE Transactions on Neural Networks* **12**(6), 674–683.
- Chen, X. and Shen, X. (1998), ‘Sieve extremum estimates for weakly dependent data’, *Econometrica* **66**(2), 298–314.
- Chen, X. and White, H. (1999), ‘Improved rates and asymptotic normality for nonparametric neural network estimators’, *IEEE Transactions on Information Theory* **45**(6), 682–691.
- Christoffersen, P. F. and Diebold, F. X. (2006), ‘Financial asset returns, direction-of-change forecasting, and volatility dynamics’, *Management Science* **52**(8), 1273–1287.
- Cybenko, G. (1989), ‘Approximation by superpositions of a sigmoidal function’, *Mathematics of Control, Signals and Systems* **2**, 303–314.

- Du, X., Fan, Y., Lv, J., Sun, T. and Vossler, P. (2021), Dimension-free average treatment effect inference with deep neural networks. Available at <https://doi.org/10.48550/arXiv.2112.01574>.
- Dubey, S. R., Singh, S. K. and Chaudhuri, B. B. (2022), ‘Activation functions in deep learning: A comprehensive survey and benchmark’, *Neurocomputing* **503**, 92–108.
- El-Sharif, I., Brown, D., Burton, B., Nixon, B. and Russell, A. (2005), ‘Evidence on the nature and extent of the relationship between oil prices and equity values in the UK’, *Energy Economics* **27**(6), 819–830.
- Engle, R. F., Ledoit, O. and Wolf, M. (2019), ‘Large dynamic covariance matrices’, *Journal of Business & Economic Statistics* **37**(2), 363–375.
- Fan, J. and Gijbels, I. (1996), *Local Polynomial Modelling and its Applications*, Chapman & Hall/CRC.
- Fan, J. and Gu, Y. (2022), Factor augmented sparse throughput deep relu neural networks for high dimensional regression. Available at <https://doi.org/10.48550/arXiv.2210.02002>.
- Fan, J., Ma, C. and Zhong, Y. (2021), ‘A selective overview of deep learning’, *Statistical Science* **36**(2), 264–290.
- Fan, J. and Yao, Q. (2003), *Nonlinear Time Series: Nonparametric and Parametric Methods*, Springer-Verlag.
- Farrell, M. H., Liang, T. and Misra, S. (2021), ‘Deep neural networks for estimation and inference’, *Econometrica* **89**(1), 181–213.
- Gao, J. (2007), *Nonlinear Time Series: Semiparametric and Nonparametric Methods*, Chapman and Hall.
- Gu, S., Kelly, B. and Xiu, D. (2020), ‘Empirical asset pricing via machine learning’, *The Review of Financial Studies* **33**(5), 2223–2273.
- Gu, S., Kelly, B. and Xiu, D. (2021), ‘Autoencoder asset pricing models’, *Journal of Econometrics* **222**(1, Part B), 429–450.
- Günther, F. and Fritsch, S. (2010), ‘Neuralnet: Training of neural networks’, *R Journal* **2**, 30–38.
- Hill, T., O’Connor, M. and Remus, W. (1996), ‘Neural network models for time series forecasts’, *Management Science* **42**(7), 1082–1092.
- Huang, J., Horowitz, J. L. and Ma, S. (2008), ‘Asymptotic properties of bridge estimators in sparse high-dimensional regression models’, *The Annals of Statistics* **36**(2), 587–613.
- Jin, Y. and Jorion, P. (1996), ‘Firm value and hedging: Evidence from U.S. oil and gas producers’, *The Journal of Finance* **61**(2), 893–919.
- Kohler, M. and Krzyżák, A. (2017), ‘Nonparametric regression based on hierarchical interaction models’, *IEEE Transactions on Information Theory* **63**(3), 1620–1630.
- Li, D., Tjøstheim, D. and Gao, J. (2016), ‘Estimation in nonlinear regression with Harris recurrent Markov chains’, *The Annals of Statistics* **44**(5), 1957–1987.
- Li, Q. and Racine, J. (2007), *Nonparametric Econometrics Theory and Practice*, Princeton University Press, New Jersey.
- Magnus, J. R. and Neudecker, H. (2007), *Matrix Differential Calculus with Applications in Statistics and Econometrics*, third edn, John Wiley & Sons Ltd.
- Minai, A. A. and Williams, R. D. (1993), ‘On the derivatives of the sigmoid’, *Neural Networks* **6**(6), 845–853.
- Mudelsee, M. (2019), ‘Trend analysis of climate time series: A review of methods’, *Earth-Science Reviews* **190**, 310–322.
- Murata, N., Yoshizawa, S. and Amari, S. (1994), ‘Network information criterion-determining the number of hidden units for an artificial neural network model’, *IEEE Transactions on Neural Networks* **5**(6), 865–872.
- Newey, W. K. and McFadden, D. (1994), Chapter 36 large sample estimation and hypothesis testing, Vol. 4 of *Handbook of Econometrics*, Elsevier, pp. 2111–2245.

- Newey, W. K. and Powell, J. L. (2003), ‘Instrumental variable estimation of nonparametric models’, *Econometrica* **71**(5), 1565–1578.
- Patrick, K. and Andres, S. (2012), ‘A score based approach to wild bootstrap inference’, *Journal of Econometric Methods* **1**(1), 23–41.
- Patton, A. J. and Sheppard, K. (2015), ‘Good volatility, bad volatility: Signed jumps and the persistence of volatility’, *The Review of Economics and Statistics* **97**(3), 683–697.
- Sauer, T. (2006), Polynomial interpolation in several variables: Lattices, differences, and ideals, *in* K. Jetter, M. D. Buhmann, W. Haussmann, R. Schaback and J. Stöckler, eds, ‘Topics in Multivariate Approximation and Interpolation’, Vol. 12 of *Studies in Computational Mathematics*, pp. 191–230.
- Schmidt-Hieber, J. (2020), ‘Nonparametric regression using deep neural networks with ReLU activation function’, *The Annals of Statistics* **48**(4), 1875–1897.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the Lasso’, *Journal of the Royal Statistical Society. Series B (Methodological)* **58**(1), 267–288.
- Wang, H. and Lin, W. (2021), Harmless overparametrization in two-layer neural networks. Available at <https://doi.org/10.48550/arXiv.2106.04795>.
- Zhong, Q., Mueller, J. and Wang, J.-L. (2022), ‘Deep learning for the partially linear Cox model’, *The Annals of Statistics* **50**(3), 1348–1375.

Online Supplementary Appendix to “A Localized Neural Network with Dependent Data: Estimation and Inference”

JITI GAO[†], BIN PENG[†] AND YANRONG YANG^{*}

[†]Monash University and ^{*}Australian National University

In this file, we first provide extra information about Sigmoidal squasher in Appendix A.1. Appendix A.2 provides some extra simulation results. In Appendix A.3, we explain how to relax the restriction about the compact set to allow for $\forall \mathbf{x}_0 \in \mathbb{R}^d$. We then introduce some notations and preliminary lemmas in Appendix A.4; their proofs are given in Appendix A.5. It is noted that Lemmas A.3 and A.4 below are of general and independent interest, as they establish useful asymptotic properties for the point-estimators of the unknown parameters of interest.

Appendix A

A.1 Derivatives of the Sigmoidal Squasher

As we adopt Sigmoidal squasher (i.e., $\sigma(x) = \frac{1}{1+\exp(-x)}$) throughout the numerical studies, we provide some of its properties below, which are useful for numerical implementation. By Minai and Williams (1993), we have

$$\sigma^{(n)}(x) = \sum_{k=1}^{n+1} (-1)^{k+1} (k-1)! S(n+1, k) \cdot \sigma(x)^k,$$

in which

$$S(n+2, k) = kS(n+1, k) + S(n+1, k-1). \tag{A.1}$$

is the recursion relation for Stirling numbers of the second kind.

Sigmoidal squasher is easy to use in the sense that we can arbitrarily choose u_σ of Lemma 3.1. Without loss of generality, we let $u_\sigma \in \{-0.5, 0.5\}$ in the numerical studies. In Figure A.1, we plot $\sigma^{(k)}(x)$ for $k = 0, \dots, 5$ for the purpose of demonstration:

A.2 Extra Simulation Results

In this section, we increase the value of d to see whether the LNN based estimation method still has good finite sample performance. The following exercises are in line with the same spirit of those provided in the supplementary file of Farrell et al. (2021), where they show, even for the i.i.d. data, one needs around 10,000 observations to generate some meaningful results with large d . We focus




Figure A.1: Plots of Sigmoidal Squasher with Its Derivatives

on Model (1.1) only. Note that when d increases, we can no longer evaluate the pre-specified points by following (4.3), as the computational overhead increases exponentially. In fact, even for the case $d = 5$, we cannot finish the computation within a few weeks using the cluster facility of Monash University if we stick to (4.3). Therefore, we consider the following points:

$$\mathbf{x}_{L,i} = -2.5 + (i - 1) \cdot 0.21_d \quad \text{with } i = 1, \dots, 26. \quad (\text{A.2})$$

As explained in Section 4, our focus is still the coverage rate. Also, for the purpose of comparison, we also consider the kernel regression with Uniform and Epanechnikov kernels respectively, and the confidence intervals are constructed using a bootstrap procedure similar to that in Section 3.2, so we omit the details. The rest settings are the same as in Section 4.

As shown in Table A.1, the LNN approach shows much better performance when $d = 5$, and the coverage rates move towards to 95% much faster. For the kernel methods, the coverage rates increase extremely slow for Epanechnikov kernel, and does not seem to increase at all for Uniform kernel. Note that the points selected in (A.2) are not on the boundary, so the results should not suffer boundary effects. It is not our intention to criticise the traditional kernel method, but LNN does show some superiority in terms of finite sample study.

Due to limits of our computational power, we no longer further increase the value of d , and also acknowledge the fact our code may not be as efficient as in some well established statistical packages.

Table A.1: Coverage rates ($d = 5$)

T	NN ($u_\sigma = -0.5$)	NN ($u_\sigma = 0.5$)	Epanechnikov	Uniform
2400	0.658	0.672	0.641	0.726
4800	0.861	0.848	0.682	0.722
9600	0.899	0.889	0.707	0.718

A.3 From $[-a, a]^d$ to \mathbb{R}^d

Consider Model (1.1) only for simplicity. We provide two treatments to relax the restriction on a .

Treatment 1: Note that in Section 3, we require $g(\mathbf{x})$ to be defined on a compact set, but do not impose restriction on the range of $\{\mathbf{x}_t\}$. In fact, for time series data, it may make more sense to assume that a is diverging, which is indeed achievable. Suppose that \mathbf{x}_t follows a sub-Gaussian distribution. In this case, we still need to require $\sqrt{Th^d}\hat{\sigma}_1^{-1} \rightarrow \infty$ in view of Lemma A.3 below, which after some algebra is equivalent to asking for

$$\sqrt{\log(Th^d)} \cdot a \rightarrow \infty.$$

Apparently, there is a price that we have to pay, which is the slow rate of convergence. A similar treatment has also been discussed in Li et al. (2016) for example, so we do not further elaborate it here.

Treatment 2: Alternatively, we can modify the construction of LNN from a nonparametric viewpoint. In the literature of kernel regression, one normally pre-specifies a point of interest (e.g., \mathbf{x}_0), and investigates a small area nearby only which is usually decided by some bandwidth(s) converging to 0. As a result, the parameters obtained from the estimation procedure usually vary with respect to the point of interest. In other words, when evaluating different points from the test set, the number of parameters to be estimated will be proportional to the cardinality of the training set (although estimation is always carried on using the same training set). Provided a large test set, it may create lots of overhead from a computational viewpoint, but the advantage is that in theory it allows us to consider $\forall \mathbf{x}_0 \in \mathbb{R}^d$.

That said, consider Model (1.1) as an example. For $\forall \mathbf{x}_0 \in \mathbb{R}^d$, we consider the following objective function:

$$Q_L(\boldsymbol{\theta}) = \sum_{t=1}^T [y_t - s(\mathbf{x}_t | \mathbf{x}_0, \boldsymbol{\theta})]^2 \cdot I_{0,h}(\mathbf{x}_t), \quad (\text{A.3})$$

where the subscript L infers the local version, $\boldsymbol{\theta}$ is $d_q \times 1$ vector satisfying $\|\boldsymbol{\theta}\| < \infty$, we let $I_{0,h}(\mathbf{x}_t) := I(\mathbf{x}_t \in C_{\mathbf{x}_0,h})$ for short, and $C_{\mathbf{x}_0,h}$ is defined in (3.3) already. Then the OLS estimate of $\tilde{\boldsymbol{\lambda}}$ defined in Lemma 3.2 is obtained by

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} Q_L(\boldsymbol{\theta})$$

and, accordingly, the estimate of $g(\mathbf{x}_0)$ is defined by

$$\hat{g}_L(\mathbf{x}_0) = s(\mathbf{x}_0 | \mathbf{x}_0, \hat{\boldsymbol{\theta}}).$$

Note that

$$\begin{aligned} \hat{g}_L(\mathbf{x}_0) - g(\mathbf{x}_0) &\simeq s(\mathbf{x}_0 | \mathbf{x}_0, \hat{\boldsymbol{\theta}}) - s(\mathbf{x}_0 | \mathbf{x}_0, \tilde{\boldsymbol{\lambda}}) \\ &\simeq \mathbf{m}(\mathbf{x}_0 | \mathbf{x}_0)^\top \mathbf{H} \mathbf{H}^{-1} \mathbf{D}^{\top,-1} (\hat{\boldsymbol{\theta}} - \tilde{\boldsymbol{\lambda}}), \end{aligned} \quad (\text{A.4})$$

where the first step is due to Lemma A.1, and the last step follows from (A.14) of the online supplementary appendix. Here, it should be understood that $\tilde{\boldsymbol{\lambda}}$ is decided by \mathbf{x}_0 only.

We summarize the results in the following corollary.

Corollary A.1. *Suppose that Assumptions 1 and 2 hold. As $(1/h, Th^d) \rightarrow (\infty, \infty)$, for $\forall \mathbf{x}_0 \in \mathbb{R}^d$,*

1. $\sqrt{Th^d} \mathbf{H}^{-1} (\mathbf{D}^\top)^{-1} (\hat{\boldsymbol{\theta}} - \tilde{\boldsymbol{\lambda}} + O_P(h^p)) \rightarrow_D N(\mathbf{0}, \sigma_\varepsilon^2 \boldsymbol{\Sigma}_{\mathbf{x}_0}^{-1}),$
2. $\sqrt{Th^d} \hat{\sigma}_{\mathbf{x}_0}^{-1} (\hat{g}_L(\mathbf{x}_0) - g(\mathbf{x}_0) + O_P(h^p)) \rightarrow_D N(0, 1),$

where $\tilde{\boldsymbol{\lambda}}$ is uniquely determined by $g(\mathbf{x}_0)$.

The first result of Corollary A.1 in fact provides CLT for the LNN parameters. While some CLTs are established previously, the focus has been on the unknown function without much inference on the effective parameters (e.g., Gu et al., 2021; Farrell et al., 2021; Chen et al., 2022).

A.4 Preliminary Lemmas

For $\mathbf{a} = (a_0, a_1, \dots, a_d)^\top \in \mathbb{R}^{d+1}$, $\mathbf{r} = (r_0, r_1, \dots, r_d)^\top \in \mathbb{N}_0^{d+1}$, and $\mathbf{x} \in \mathbb{R}^d$, we let

$$f_{\mathbf{a}}(\mathbf{x}) = [(1, \mathbf{x}^\top) \mathbf{a}]^q = \sum_{|\mathbf{r}|=q} \binom{q}{\mathbf{r}} \cdot a_0^{r_0} \prod_{k=1}^d a_k^{r_k} x_k^{r_k}.$$

Obviously, we have $f_{\mathbf{a}} \in \mathcal{P}_q$, where \mathcal{P}_q is defined in (3.1).

Lemma A.1. *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a (p, \mathcal{C}) -smooth function. For $\forall \mathbf{x}_0 \in \mathbb{R}^d$, let*

$$p_q(\mathbf{x} | \mathbf{x}_0) = \sum_{0 \leq |\mathbf{J}| \leq q} \frac{1}{\mathbf{J}!} \cdot \frac{\partial^{|\mathbf{J}|} f(\mathbf{x}_0)}{\partial \mathbf{x}^{\mathbf{J}}} (\mathbf{x} - \mathbf{x}_0)^{\mathbf{J}}.$$

Then

$$\|f(\mathbf{x}) - p_q(\mathbf{x} | \mathbf{x}_0)\|_\infty \leq O(1) \|\mathbf{x} - \mathbf{x}_0\|^p,$$

where $p = q + s$, $O(1)$ depends on d and q only.

Lemma A.2. *For almost all $\mathbf{a}_1, \dots, \mathbf{a}_{d_q} \in \mathbb{R}^{d+1}$ (with respect to the Lebesgue measure in $\mathbb{R}^{(d+1) \times d_q}$), we have that $\{f_{\mathbf{a}_1}(\mathbf{x}), \dots, f_{\mathbf{a}_{d_q}}(\mathbf{x})\}$ is a basis of the linear vector space \mathcal{P}_q .*

Lemma A.3 (On Model (1.1)). *Suppose Assumptions 1 and 2 hold. As $(1/h, Th^d) \rightarrow (\infty, \infty)$, for each $\mathbf{i} \in [M]^d$,*

$$\sigma_\varepsilon^{-1} \boldsymbol{\Sigma}_{\mathbf{i}}^{1/2} \sqrt{Th^d} \mathbf{H}^{-1} \mathbf{D}^{\top, -1} (\hat{\boldsymbol{\theta}}_{\mathbf{i}} - \tilde{\boldsymbol{\lambda}}_{\mathbf{i}} + O_P(h^p)) \rightarrow_D N(\mathbf{0}, \mathbf{I}),$$

where $\boldsymbol{\Sigma}_{\mathbf{i}} = f_{\mathbf{x}}(\tilde{\mathbf{x}}_{\mathbf{i}}) \int_{[-1, 1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}$.

Lemma A.4 (On Model (1.2)). *Under Assumptions 1 and 2, for each $\mathbf{i} \in [M]^d$,*

1. $\frac{1}{T} \sum_{t=1}^T I_{i,h}(\mathbf{x}_t)[g(\mathbf{x}_t) - s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)]^2 = o_P(1)$,
2. $\left\| \frac{1}{T} \mathbf{H} \mathbf{D} \frac{\partial^2 \log L(\tilde{s}(\cdot | \hat{\boldsymbol{\Theta}}))}{\partial \boldsymbol{\theta}_i \partial \boldsymbol{\theta}_i^\top} \mathbf{D}^\top \mathbf{H} - \tilde{\boldsymbol{\Sigma}}_i \right\| = o_P(1)$,
3. $\left\| \frac{1}{T} \mathbf{H} \mathbf{D} \frac{\partial^2 \log L(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i^*))}{\partial \boldsymbol{\theta}_i \partial \boldsymbol{\theta}_i^\top} \mathbf{D}^\top \mathbf{H} - \tilde{\boldsymbol{\Sigma}}_i \right\| = o_P(1)$,
4. $\tilde{\boldsymbol{\Sigma}}_i^{-1/2} \frac{1}{\sqrt{Th^d}} \mathbf{H} \mathbf{D} \frac{\partial \log L(g(\cdot))}{\partial \boldsymbol{\theta}_i} \rightarrow_D N(\mathbf{0}, \mathbf{I})$,

where $\boldsymbol{\theta}_i^*$ lies between $\hat{\boldsymbol{\theta}}_i$ and $\tilde{\boldsymbol{\lambda}}_i$, and

$$\tilde{\boldsymbol{\Sigma}}_i = \frac{f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \phi_\varepsilon(g(\tilde{\mathbf{x}}_i))^2}{[1 - \Phi_\varepsilon(g(\tilde{\mathbf{x}}_i))] \Phi_\varepsilon(g(\tilde{\mathbf{x}}_i))} \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}.$$

A.5 Proofs

Proof of Lemma A.1:

This is Lemma 8 of Bauer and Kohler (2019), so the derivation is omitted. ■

Proof of Lemma A.2:

In what follows, let

$$\mathbf{a}_k = (a_{k,0}, a_{k,1}, \dots, a_{k,d})^\top \text{ for } k \in [d_q], \quad \text{and} \quad \tilde{\mathbf{r}} = \{\mathbf{r} \in \mathbb{N}_0^{d+1} \mid |\mathbf{r}| = q\}.$$

It suffices to show that $f_{\mathbf{a}_1}(\mathbf{x}), \dots, f_{\mathbf{a}_{d_q}}(\mathbf{x})$ are linearly independent. To do this, let $b_1, \dots, b_{d_q} \in \mathbb{R}$ be such that

$$\sum_{k=1}^{d_q} b_k f_{\mathbf{a}_k}(\mathbf{x}) = 0. \tag{A.5}$$

As we explained under (3.1), the monomials involved in (3.1) are linearly independent. Thus, (A.5) implies that

$$\sum_{k=1}^{d_q} b_k \prod_{j=0}^d a_{k,j}^{r_j} = 0 \quad \text{for } \forall \mathbf{r} \in \tilde{\mathbf{r}}.$$

Note that $\#\tilde{\mathbf{r}} = d_q$ by design, so we can construct a one-to-one relationship between $k \in [d_q]$ and $\mathbf{r} \in \tilde{\mathbf{r}}$. Using this relationship, we can construct $p(\mathbf{x}) \in \mathcal{P}_q$ as follows:

$$p(\mathbf{x}) = \sum_{k=1}^{d_q} b_k \prod_{j=1}^d x_j^{r_j} = \sum_{\mathbf{r} \in \tilde{\mathbf{r}}} b_{\mathbf{r}} \prod_{j=1}^d x_j^{r_j},$$

which satisfies

$$p(\mathbf{a}_k) = 0 \quad \text{for } \forall k \in [d_q]. \tag{A.6}$$

Position 4 in Sauer (2006) implies that (A.6) has the only solution $p(\cdot) = 0$ in \mathcal{P}_q for Lebesgue almost all $\mathbf{a}_1, \dots, \mathbf{a}_{d_q} \in \mathbb{R}^{d+1}$, which in turn implies $b_1 = \dots = b_{d_q} = 0$. The proof is now completed. \blacksquare

Proof of Lemma 3.1:

Before proceeding, we would like to point out that in what follows, \mathbf{C} is a constant and is for the purpose of rescaling only.

By Assumption 1.2, there is a point $u_\sigma \in \mathbb{R}$ such that none of the derivatives up to the order q is 0 at u_σ . Thus, we construct the following one-layer NN:

$$\begin{aligned}
& \sum_{k=1}^{q+1} (-1)^{q+k-1} \cdot \frac{\mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \binom{q}{k-1} \cdot \sigma \left(\frac{k-1}{\mathbf{C}} \cdot (x-x_0) + u_\sigma \right) \\
&= \sum_{k=0}^q (-1)^{q+k} \cdot \frac{\mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \binom{q}{k} \cdot \sigma \left(\frac{k}{\mathbf{C}} \cdot (x-x_0) + u_\sigma \right) \\
&= (-1)^q \frac{\mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \sum_{k=0}^q (-1)^k \cdot \binom{q}{k} \cdot \sigma \left(\frac{k}{\mathbf{C}} \cdot (x-x_0) + u_\sigma \right), \tag{A.7}
\end{aligned}$$

in which the definitions of γ_k 's and β_k 's are obvious.

By Assumption 1.2 again, $\sigma(\cdot)$ is $q+1$ times continuously differentiable. Thus, it can be expanded in a Taylor series with Lagrange remainder around u_σ up to order q :

$$\begin{aligned}
& \sum_{k=0}^q (-1)^k \cdot \binom{q}{k} \cdot \sigma \left(\frac{k}{\mathbf{C}} \cdot (x-x_0) + u_\sigma \right) \\
&= \sum_{k=0}^q (-1)^k \cdot \binom{q}{k} \cdot \left(\sum_{j=0}^q \frac{\sigma^{(j)}(u_\sigma) \cdot ((x-x_0)k)^j}{\mathbf{C}^j \cdot j!} + \frac{\sigma^{(q+1)}(\xi_k) \cdot ((x-x_0)k)^{q+1}}{\mathbf{C}^{q+1} \cdot (q+1)!} \right) \\
&= \sum_{j=0}^q \frac{\sigma^{(j)}(u_\sigma) \cdot (x-x_0)^j}{\mathbf{C}^j \cdot j!} \sum_{k=0}^q (-1)^k \cdot k^j \cdot \binom{q}{k} \\
& \quad + \frac{(x-x_0)^{q+1}}{\mathbf{C}^{q+1} \cdot (q+1)!} \sum_{k=0}^q (-1)^k \cdot k^{q+1} \cdot \sigma^{(q+1)}(\xi_k) \cdot \binom{q}{k}, \tag{A.8}
\end{aligned}$$

where $\xi_k \in [u_\sigma - \frac{k}{\mathbf{C}} \cdot |x-x_0|, u_\sigma + \frac{k}{\mathbf{C}} \cdot |x-x_0|]$ for all $0 \leq k \leq q$.

Note that

$$\begin{aligned}
\sum_{k=0}^q (-1)^k \cdot k^j \cdot \binom{q}{k} &= q! (-1)^q \cdot \frac{1}{q!} \sum_{k=0}^q (-1)^{k-q} \cdot (q-(q-k))^j \cdot \binom{q}{q-k} \\
&= q! (-1)^q \cdot \frac{1}{q!} \sum_{k=0}^q (-1)^{q-k} \cdot (q-(q-k))^j \cdot \binom{q}{q-k} \\
&= q! (-1)^q \cdot \left\{ \begin{matrix} j \\ q \end{matrix} \right\},
\end{aligned}$$

where $\left\{ \begin{matrix} j \\ q \end{matrix} \right\}$ is the Stirling number of the second kind. The Stirling number of the second kind describes the number of options to split a set of j elements into n non-empty subsets, which is equal

to 0 for $0 \leq j < n$, and is equal to 1 for $j = n$. The result holds true for all $j, n \in \mathbb{N}$ (Abramovitz and Stegun, 1972, p. 825).

Thus, we can further simplify the right hand side of (A.8), and write

$$\begin{aligned} & \sum_{k=0}^q (-1)^k \cdot \binom{q}{k} \cdot \sigma \left(\frac{k}{\mathbf{C}} \cdot (x - x_0) + u_\sigma \right) \\ &= \frac{\sigma^{(q)}(u_\sigma) \cdot (x - x_0)^q}{\mathbf{C}^q} \cdot (-1)^q \\ & \quad + \frac{(x - x_0)^{q+1}}{\mathbf{C}^{q+1} \cdot (q+1)!} \sum_{k=0}^q (-1)^k \cdot k^{q+1} \cdot \sigma^{(q+1)}(\xi_k) \cdot \binom{q}{k}, \end{aligned}$$

which in connection with (A.7) yields that

$$\begin{aligned} & (-1)^q \frac{\mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \sum_{k=0}^q (-1)^k \cdot \binom{q}{k} \cdot \sigma \left(\frac{k}{\mathbf{C}} \cdot (x - x_0) + u_\sigma \right) \\ &= (-1)^q \frac{\mathbf{C}^q}{\sigma^{(q)}(u_\sigma)} \left\{ \frac{\sigma^{(q)}(u_\sigma) \cdot (x - x_0)^q}{\mathbf{C}^q} \cdot (-1)^q \right. \\ & \quad \left. + \frac{(x - x_0)^{q+1}}{\mathbf{C}^{q+1} \cdot (q+1)!} \sum_{k=0}^q (-1)^k \cdot k^{q+1} \cdot \sigma^{(q+1)}(\xi_k) \cdot \binom{q}{k} \right\} \\ &= (x - x_0)^q + \frac{(-1)^q (x - x_0)^{q+1}}{\mathbf{C} \cdot \sigma^{(q)}(u_\sigma) \cdot (q+1)!} \sum_{k=0}^q (-1)^k \cdot k^{q+1} \cdot \sigma^{(q+1)}(\xi_k) \cdot \binom{q}{k}. \end{aligned}$$

In view of Assumption 1.2 and \mathbf{C} being a fixed value, the proof is now completed. ■

Proof of Lemma 3.2:

By Lemma A.2, we can reconstruct all of $\{m_i(\mathbf{x} | \mathbf{x}_0)\}$ as follows:

$$\mathbf{m}(\mathbf{x} | \mathbf{x}_0) = \mathbf{D}\mathbf{A}(\mathbf{x} | \mathbf{x}_0)$$

where

$$\begin{aligned} \mathbf{D} &= \{d_{ij}\}_{d_q \times d_q} = (\mathbf{d}_1, \dots, \mathbf{d}_{d_q}), \\ \mathbf{A}(\mathbf{x} | \mathbf{x}_0) &= \left([(1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_1]^q, \dots, [(1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_{d_q}]^q \right)^\top. \end{aligned}$$

Note that the rotation matrix \mathbf{D} is determined by $\boldsymbol{\alpha}_j$'s only, so they are fixed.

Apparently, we have an issue of identification here, because for example we can arbitrarily rescale $\boldsymbol{\alpha}_j$'s, and modify \mathbf{D} accordingly without changing $m_i(\mathbf{x} | \mathbf{x}_0)$ as follows:

$$\mathbf{D}\mathbf{A}(\mathbf{x} | \mathbf{x}_0) = \mathbf{D}\mathbf{B}\mathbf{B}^{-1}\mathbf{A}(\mathbf{x} | \mathbf{x}_0),$$

in which \mathbf{B} is full rank. Therefore, for the purpose of identification, we regulate $\boldsymbol{\alpha}_j$'s as follows:

$$[\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_{d_q}] = \frac{1}{d+1} \cdot \mathbf{I}_h \mathbf{W} \tag{A.9}$$

in which $\mathbf{I}_h = \text{diag}\{h, \mathbf{I}_d\}$, and \mathbf{W} is defined in the body of this lemma. As a result, for $\forall j \in [d_q]$,

$$\begin{aligned} \sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} |(1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j| &= \sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \frac{1}{d+1} |(h, \mathbf{x}^\top - \mathbf{x}_0^\top) \mathbf{w}_j| \\ &\leq \frac{h}{d+1} \sqrt{d+1} \cdot \max_j \|\mathbf{w}_j\| = h, \end{aligned}$$

so we can invoke Lemma 3.1 later on.

Treating $(1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j$ as a whole and using Lemma 3.1, we write

$$\begin{aligned} &\sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \left| m_i(\mathbf{x} | \mathbf{x}_0) - \sum_{j=1}^{d_q} d_{ij} \cdot \sum_{k=1}^{q+1} \gamma_k \cdot \sigma \left(\beta_k (1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j + u_\sigma \right) \right| \\ &\leq \sum_{j=1}^{d_q} |d_{ij}| \cdot \sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \left| [(1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j]^q - \sum_{k=1}^{q+1} \gamma_k \cdot \sigma \left(\beta_k (1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j + u_\sigma \right) \right| \\ &= O(h^{q+1}), \end{aligned} \tag{A.10}$$

where the last line follows from Lemma 3.1. Also, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and u_σ are known as discussed in Remark 3.1.

Thus, we can further write

$$\begin{aligned} &\sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \left| p(\mathbf{x} | \mathbf{x}_0, \boldsymbol{\lambda}) - \sum_{j=1}^{d_q} \boldsymbol{\lambda}^\top \mathbf{d}_j \cdot \sum_{k=1}^{q+1} \gamma_k \cdot \sigma \left(\beta_k (1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j + u_\sigma \right) \right| \\ &= \sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \left| \boldsymbol{\lambda}^\top \mathbf{m}(\mathbf{x} | \mathbf{x}_0) - \sum_{j=1}^{d_q} \boldsymbol{\lambda}^\top \mathbf{d}_j \cdot \sum_{k=1}^{q+1} \gamma_k \cdot \sigma \left(\beta_k (1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j + u_\sigma \right) \right| \\ &\leq \|\boldsymbol{\lambda}\| \cdot \sup_{\mathbf{x} \in C_{\mathbf{x}_0, h}} \left\| \mathbf{m}(\mathbf{x} | \mathbf{x}_0) - \sum_{j=1}^{d_q} \mathbf{d}_j \cdot \sum_{k=1}^{q+1} \gamma_k \cdot \sigma \left(\beta_k (1, \mathbf{x}^\top - \mathbf{x}_0^\top) \boldsymbol{\alpha}_j + u_\sigma \right) \right\| \\ &= O(h^{q+1}), \end{aligned} \tag{A.11}$$

where the last line follows from the facts that d_q is fixed and $\|\boldsymbol{\lambda}\| = O(1)$.

Finally, let $\tilde{\boldsymbol{\lambda}} = (\tilde{\lambda}_1, \dots, \tilde{\lambda}_{d_q})^\top$ with $\tilde{\lambda}_j = \boldsymbol{\lambda}^\top \mathbf{d}_j$. Further, in view of the definitions of $\boldsymbol{\sigma}(\mathbf{x} | \mathbf{x}_0)$ and $(\boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_{d_q(q+1)})$ in the body of this lemma and (A.9), the proof is then completed. \blacksquare

Proof of Lemma 3.3:

Before starting the proof, we introduce a few notations to facilitate the development. First, recall that in the body of this theorem, we have defined

$$\tilde{s}(\mathbf{x} | \tilde{\boldsymbol{\Lambda}}) = \sum_{\mathbf{i} \in [M]^d} I_{i, h}(\mathbf{x}) \cdot s(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}),$$

where $s(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}) = (\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} \otimes \boldsymbol{\gamma})^\top \boldsymbol{\sigma}(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}})$ by the definition of Lemma 3.2. Second, note that the leading terms of the q^{th} order Taylor expansion of $g(\mathbf{x})$ at each $\tilde{\mathbf{x}}_{\mathbf{i}} \in (-a, a)^d$ can be written as follows:

$$\sum_{0 \leq |\mathbf{J}| \leq q} \frac{1}{\mathbf{J}!} \cdot \frac{\partial^{|\mathbf{J}|} g(\tilde{\mathbf{x}}_{\mathbf{i}})}{\partial \mathbf{x}^{\mathbf{J}}} (\mathbf{x} - \tilde{\mathbf{x}}_{\mathbf{i}})^{\mathbf{J}} = \boldsymbol{\lambda}_{\mathbf{i}}^\top \mathbf{m}(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}) := p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}),$$

where the definition of $\boldsymbol{\lambda}_{\mathbf{i}}$ should be obvious in view of the definition of $\mathbf{m}(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}})$ according to (3.2).

We are now ready to start the proof, and write

$$\begin{aligned} g(\mathbf{x}) - \tilde{s}(\mathbf{x} | \tilde{\mathbf{\Lambda}}) &= \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}) - \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot s(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}) \\ &\quad + g(\mathbf{x}) - \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}). \end{aligned}$$

Note that by Lemma 3.2 we choose $\tilde{\boldsymbol{\lambda}}_{\mathbf{i}}$ which fulfils the relationship: $\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} = \mathbf{D}\boldsymbol{\lambda}_{\mathbf{i}}$. It is worth mentioning that although $(\tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}})$ vary with respect to \mathbf{i} , the rotation matrix \mathbf{D} in facts is solely determined by \mathbf{W} of Lemma 3.2. Therefore, without loss of generality, we can fix \mathbf{W} over \mathbf{i} , as it is user chosen. Then \mathbf{D} remains the same in view of the proof of Lemma 3.2.

Next, we write

$$\begin{aligned} &\left\| \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}) - g(\mathbf{x}) \right\|_{\infty} \\ &= \left\| \sum_{\mathbf{i} \in [M]^d} [p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}) - g(\mathbf{x})] \cdot I_{\mathbf{i},h}(\mathbf{x}) \right\|_{\infty} \\ &\leq \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot \sup_{\mathbf{x} \in C_{\mathbf{x}_{0\mathbf{i}},h}} |p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}) - g(\mathbf{x})| = O(h^p), \end{aligned} \quad (\text{A.12})$$

where the inequality follows from the definition of $I_{\mathbf{i},h}(\mathbf{x})$, and the last step follows from Lemma A.1. Also, we can obtain that

$$\begin{aligned} &\left\| \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot s(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}) - \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}}) \right\|_{\infty} \\ &\leq \sum_{\mathbf{i} \in [M]^d} I_{\mathbf{i},h}(\mathbf{x}) \cdot \sup_{\mathbf{x} \in C_{\mathbf{x}_{0\mathbf{i}},h}} |s(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}) - p(\mathbf{x} | \tilde{\mathbf{x}}_{\mathbf{i}}, \boldsymbol{\lambda}_{\mathbf{i}})| = O(h^{q+1}), \end{aligned} \quad (\text{A.13})$$

where the inequality follows from the definition of $I_{\mathbf{i},h}(\mathbf{x})$, and the last step follows from Lemma 3.2 by letting $\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} = \mathbf{D}\boldsymbol{\lambda}_{\mathbf{i}}$.

Therefore, based on (A.12) and (A.13), we obtain

$$\|g(\mathbf{x}) - \tilde{s}(\mathbf{x} | \tilde{\mathbf{\Lambda}})\|_{\infty} = O(h^p),$$

where $p = q + s$. The proof is now completed. ■

A.5.1 Model (1.1)

Proof of Lemma A.3:

First, we expand the expression of $\hat{\boldsymbol{\theta}}_{\mathbf{i}}$ as follows:

$$\hat{\boldsymbol{\theta}}_{\mathbf{i}} - \tilde{\boldsymbol{\lambda}}_{\mathbf{i}} = \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^{\top} \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} y_t - \tilde{\boldsymbol{\lambda}}_{\mathbf{i}}$$

$$\begin{aligned}
&= \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{s}(\mathbf{x}_t | \tilde{\mathbf{\Lambda}}) - \tilde{\boldsymbol{\lambda}}_i \\
&\quad + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\mathbf{\Lambda}})] \\
&\quad + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \varepsilon_t \\
&= \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\mathbf{\Lambda}})] \\
&\quad + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \varepsilon_t,
\end{aligned}$$

where the third equality follows from the fact that $I_{i,h}(\mathbf{x}_t)I_{j,h}(\mathbf{x}_t) = 0$ for $\mathbf{i} \neq \mathbf{j}$, and the definition of $\tilde{s}(\mathbf{x}_t | \tilde{\mathbf{\Lambda}})$. Below, we consider the terms on the right hand side one by one.

Recall the argument of Remark 3.5, so we immediately obtain that

$$\|\tilde{\mathbf{x}}_{i,t}^\top \mathbf{D}^\top \mathbf{H} - I_{i,h}(\mathbf{x}_t) \mathbf{m}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)^\top \cdot \mathbf{H}\| = O(h), \quad (\text{A.14})$$

where the last line follows from the definition of \mathbf{H} and $\max_j |\mathbf{n}_j| = q$. For notational simplicity, we further denote

$$I_{i,h}(\mathbf{x}_t) \mathbf{H} \cdot \mathbf{m}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) := \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) = (\tilde{m}_1(\mathbf{x}_t | \tilde{\mathbf{x}}_i), \dots, \tilde{m}_{d_q}(\mathbf{x}_t | \tilde{\mathbf{x}}_i))^\top. \quad (\text{A.15})$$

First, we consider $\frac{1}{Th^d} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top$, and write

$$\begin{aligned}
&\frac{1}{Th^d} \sum_{t=1}^T \mathbf{H} \mathbf{D} E \left[\tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right] \mathbf{D}^\top \mathbf{H} \\
&= \frac{1}{h^d} \mathbf{H} \mathbf{D} E[\tilde{\mathbf{x}}_{i,1} \tilde{\mathbf{x}}_{i,1}^\top] \mathbf{D}^\top \mathbf{H} \\
&= \frac{1}{h^d} E[I_{i,h}(\mathbf{x}_1) \mathbf{H} \cdot \mathbf{m}(\mathbf{x}_1 | \tilde{\mathbf{x}}_i) \mathbf{m}(\mathbf{x}_1 | \tilde{\mathbf{x}}_i)^\top \cdot \mathbf{H}] \cdot (1 + o(1)) \\
&= \frac{1}{h^d} \int_{\mathbf{x} \in C_{\mathbf{x}_{0i},h}} \mathbf{H} \cdot \mathbf{m}(\mathbf{x} | \tilde{\mathbf{x}}_i) \mathbf{m}(\mathbf{x} | \tilde{\mathbf{x}}_i)^\top \cdot \mathbf{H} \cdot f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x} \cdot (1 + o(1)) \\
&= f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x} \cdot (1 + o(1)) > 0
\end{aligned} \quad (\text{A.16})$$

where the first equality follows from Assumption 2.1, the second equality follows from (A.14), and the last step follows from Assumption 2.2 and the definition of $\mathbf{m}(\mathbf{x} | \mathbf{0})$ given in (3.2).

Next, we consider the second moment. Define $\mathbf{X}_i = E[\mathbf{H} \mathbf{D} \tilde{\mathbf{x}}_{t,i} \tilde{\mathbf{x}}_{t,i}^\top \mathbf{D}^\top \mathbf{H}]$ for notational simplicity, and by the development similar to those given for (A.18) below it is easy to see that

$$E \left\| \frac{1}{Th^d} \sum_{t=1}^T \mathbf{H} \mathbf{D} \tilde{\mathbf{x}}_{t,i} \tilde{\mathbf{x}}_{t,i}^\top \mathbf{D}^\top \mathbf{H} - \mathbf{X}_i \right\|^2 = o(1), \quad (\text{A.17})$$

so we omit the details for now.

Also, we note

$$\begin{aligned}
& \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\Lambda})] \\
&= \mathbf{D}^\top \mathbf{H} \left(\mathbf{H} \mathbf{D} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \mathbf{D}^\top \mathbf{H} \right)^{-1} \mathbf{H} \mathbf{D} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\Lambda})] \\
&= \mathbf{D}^\top \mathbf{H} (\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \Delta \mathbf{G},
\end{aligned}$$

where

$$\begin{aligned}
\tilde{\mathbf{X}}_i &= (\tilde{\mathbf{x}}_{i,1}, \dots, \tilde{\mathbf{x}}_{i,T})^\top, \\
\Delta \mathbf{G} &= (I_{i,h}(\mathbf{x}_1)[g(\mathbf{x}_1) - \tilde{s}(\mathbf{x}_1 | \tilde{\Lambda})], \dots, I_{i,h}(\mathbf{x}_T)[g(\mathbf{x}_T) - \tilde{s}(\mathbf{x}_T | \tilde{\Lambda})])^\top.
\end{aligned}$$

By Lemma 3.3, it is easy to see that

$$\frac{1}{Th^d} E \|\Delta \mathbf{G}\|^2 = \frac{1}{h^d} E \left[I_{i,h}(\mathbf{x}_1)[g(\mathbf{x}_1) - \tilde{s}(\mathbf{x}_1 | \tilde{\Lambda})]^2 \right] = O(h^{2p}).$$

Then we can write

$$\begin{aligned}
& \|(\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \Delta \mathbf{G}\|^2 \\
&= \Delta \mathbf{G}^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H} (\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} (\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \Delta \mathbf{G} \\
&\leq \frac{1}{Th^d} \lambda_{\max} \left\{ \left(\frac{1}{Th^d} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H} \right)^{-1} \right\} \Delta \mathbf{G}^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H} (\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \Delta \mathbf{G} \\
&\leq O_P(1) \lambda_{\max} \{ \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H} (\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \tilde{\mathbf{X}}_i \mathbf{D}^\top \mathbf{H})^{-1} \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_i^\top \} \cdot \|\Delta \mathbf{G}\|^2 / (Th^d) \\
&= O_P(h^{2p}),
\end{aligned}$$

where the first inequality follows from the exercise 5 on page 267 of Magnus and Neudecker (2007), and the second inequality follows from (A.16) and (A.17).

Based on the above development, we can conclude that

$$\left\| \mathbf{H}^{-1} \mathbf{D}^{\top, -1} \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} \tilde{\mathbf{x}}_{i,t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{i,t} [g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\Lambda})] \right\| = O_P(h^p).$$

Finally, in order to establish the asymptotic distribution, we just need to focus on $\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t$ in view of (A.14). Write

$$\begin{aligned}
& E \left[\left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t \right) \left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t \right)^\top \right] \\
&= \frac{1}{Th^d} \sum_{t=1}^T \sum_{s=1}^T E[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_s | \tilde{\mathbf{x}}_i)^\top \varepsilon_t \varepsilon_s] \\
&= \frac{1}{Th^d} \sum_{t=1}^T \sigma_\varepsilon^2 E[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)^\top]
\end{aligned}$$

$$\begin{aligned}
& + \frac{1}{h^d} \sum_{t=1}^{T-1} (1-t/T) E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}] \\
& + \frac{1}{h^d} \sum_{t=1}^{T-1} (1-t/T) E[\tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}], \tag{A.18}
\end{aligned}$$

where the last two terms are the same up to a transpose operation.

The term $\frac{1}{h^d} \sum_{t=1}^{T-1} (1-t/T) E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}]$ on the right hand side can be bounded as follows.

$$\begin{aligned}
& \|E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}]\| \\
& \leq O(1) \alpha(t)^{\nu/(2+\nu)} \{E\|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \varepsilon_1\|^{2+\nu}\}^{2/(2+\nu)} \\
& \leq O(1) \alpha(t)^{\nu/(2+\nu)} \left\{E\|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \varepsilon_1\|^{2+\nu} \frac{1}{h^d}\right\}^{2/(2+\nu)} \cdot (h^d)^{2/(2+\nu)} \\
& = O((h^d)^{2/(2+\nu)}) \alpha(t)^{\nu/(2+\nu)}.
\end{aligned}$$

Thus, we have

$$\begin{aligned}
& \left\| \frac{1}{h^d} \sum_{t=1}^{T-1} (1-t/T) E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}] \right\| \\
& \leq O(1) \frac{1}{h^d} \sum_{t=1}^{d_T} \left\| E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}] \right\| \\
& \quad + O(1) \frac{1}{h^d} \sum_{t=d_T+1}^T \left\| E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}] \right\| \\
& \leq O(1) h^d \sum_{t=1}^{d_T} E \left\| \frac{1}{h^d} I_{i,h}(\mathbf{x}_1) \frac{1}{h^d} I_{i,h}(\mathbf{x}_{1+t}) \varepsilon_1 \varepsilon_{1+t} \right\| \\
& \quad + O(1) \frac{1}{h^d} \sum_{t=d_T+1}^T \left\| E[\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1) \tilde{\mathbf{m}}(\mathbf{x}_{1+t} | \tilde{\mathbf{x}}_1)^\top \varepsilon_1 \varepsilon_{1+t}] \right\| \\
& \leq O(1) h^d d_T + O(1) \frac{(h^d)^{2/(2+\nu)}}{h^d} \sum_{t=d_T+1}^T \alpha^{\nu/(2+\nu)}(t) \\
& = O(1) h^d d_T + O(1) \frac{1}{h^{\frac{\nu d}{2+\nu}}} \sum_{t=d_T+1}^T \alpha^{\nu/(2+\nu)}(t) = o(1), \tag{A.19}
\end{aligned}$$

where $\nu > 0$ is defined in Assumption 2.1, and the last step follows from Assumption 2.1 that we can choose d_T to ensure

$$d_T h^d \rightarrow 0 \quad \text{and} \quad \sum_{t=d_T+1}^T \alpha^{\nu/(2+\nu)}(t) = o\left(h^{\frac{\nu d}{2+\nu}}\right), \tag{A.20}$$

which can be achieved by choosing $d_T = \lfloor T^{c_1} \rfloor$, $h = T^{-c_2}$ and $\alpha(t) = t^{-c_3}$ for some suitable $c_1 > 0$, $c_2 > 0$ and $c_3 > 0$, for example.

Thus, we can conclude that

$$\begin{aligned}
& E \left[\left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t \right) \left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t \right)^\top \right] \\
&= \frac{1}{Th^d} \sum_{t=1}^T \sigma_\varepsilon^2 E[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)^\top] + o(1) \\
&\rightarrow \sigma_\varepsilon^2 f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}. \tag{A.21}
\end{aligned}$$

Below, we further use small-block and large-block to prove the normality. To employ the small-block and large-block arguments, we partition the set $\{1, \dots, T\}$ into $2k_T + 1$ subsets with large blocks of size l_T and small blocks of size s_T and the last remaining set of size $T - k_T(l_T + s_T)$, where l_T and s_T are selected such that

$$s_T \rightarrow \infty, \quad \frac{s_T}{l_T} \rightarrow 0, \quad \frac{l_T^{1+\nu}}{(Th^d)^{\frac{\nu}{2}}} \rightarrow 0, \quad \text{and} \quad k_T \equiv \left\lfloor \frac{T}{l_T + s_T} \right\rfloor,$$

and ν is defined in Assumption 2.1.

For $j = 1, \dots, k_T$, define

$$\begin{aligned}
\boldsymbol{\xi}_j &= \sum_{t=(j-1)(l_T+s_T)+1}^{jl_T+(j-1)s_T} \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t, & \boldsymbol{\eta}_j &= \sum_{t=jl_T+(j-1)s_T+1}^{j(l_T+s_T)} \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t, \\
\boldsymbol{\zeta} &= \sum_{t=k_T(l_T+s_T)+1}^T \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t.
\end{aligned}$$

Note that $\alpha(T) = o(1/T)$ and $k_T s_T / T \rightarrow 0$. By direct calculation, we immediately obtain that

$$\frac{1}{T} E \left\| \sum_{j=1}^{k_T} \boldsymbol{\eta}_j \right\|^2 \rightarrow 0 \quad \text{and} \quad \frac{1}{T} E \|\boldsymbol{\zeta}\|^2 \rightarrow 0.$$

Therefore,

$$\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \varepsilon_t = \frac{1}{\sqrt{T}} \sum_{j=1}^{k_T} \boldsymbol{\xi}_j + o_P(1).$$

By Proposition 2.6 of Fan and Yao (2003), we have as $T \rightarrow 0$

$$\begin{aligned}
& \left| E \left[\exp \left(\frac{iw}{\sqrt{T}} \sum_{j=1}^{k_T} \boldsymbol{\xi}_j \right) \right] - \prod_{j=1}^{k_T} E \left[\exp \left(\frac{iw \boldsymbol{\xi}_j}{\sqrt{T}} \right) \right] \right| \\
&\leq 16(k_T - 1) \alpha(s_T) \rightarrow 0,
\end{aligned}$$

where i is the imaginary unit. In connection with (A.18)-(A.21), the Feller condition is fulfilled as follows:

$$\frac{1}{T} \sum_{j=1}^{k_T} E[\boldsymbol{\xi}_j \boldsymbol{\xi}_j^\top] \rightarrow \sigma_\varepsilon^2 f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}.$$

Also, we note that

$$\begin{aligned}
E[\|\boldsymbol{\xi}_1\|^2 \cdot I(\|\boldsymbol{\xi}_1\| \geq \epsilon\sqrt{T})] &\leq \{E\|\boldsymbol{\xi}_1\|^{2+\nu}\}^{\frac{2}{2+\nu}} \left\{E[I(\|\boldsymbol{\xi}_1\| \geq \epsilon\sqrt{T})]\right\}^{\frac{\nu}{2+\nu}} \\
&\leq \{E\|\boldsymbol{\xi}_1\|^{2+\nu}\}^{\frac{2}{2+\nu}} \left\{\frac{E\|\boldsymbol{\xi}_1\|^{2+\nu}}{\epsilon^{2+\nu}T^{\frac{2+\nu}{2}}}\right\}^{\frac{\nu}{2+\nu}} \\
&= \frac{1}{\epsilon^\nu T^{\frac{\nu}{2}}} \{E\|\boldsymbol{\xi}_1\|^{2+\nu}\}^{\frac{1}{2+\nu} \cdot (2+\nu)} \\
&= O(1) \frac{l_T^{2+\nu}}{\epsilon^\nu T^{\frac{\nu}{2}}} E\|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1)\varepsilon_1\|^{2+\nu} \cdot \frac{1}{h^{\frac{d}{2} \cdot (2+\nu)}} \\
&= O(1) \frac{l_T^{2+\nu}}{\epsilon^\nu (Th^d)^{\frac{\nu}{2}}} \cdot \frac{1}{h^d} E\|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_1)\varepsilon_1\|^{2+\nu} \\
&= O(1) \frac{l_T^{2+\nu}}{(Th^d)^{\frac{\nu}{2}}},
\end{aligned}$$

where the first inequality follows from Hölder inequality, the second inequality follows from Chebyshev's inequality, and the second equality follows from Minkowski inequality. Consequently,

$$\frac{1}{T} \sum_{j=1}^{k_T} E[\|\boldsymbol{\xi}_j\|^2 \cdot I(\|\boldsymbol{\xi}_j\| \geq \epsilon\sqrt{T})] = O\left(\frac{k_T l_T^{2+\nu}}{T(Th^d)^{\frac{\nu}{2}}}\right) = O\left(\frac{l_T^{1+\nu}}{(Th^d)^{\frac{\nu}{2}}}\right) = o(1),$$

where the last step follows from the choice of l_T as specified above. Therefore, the Lindberg condition is justified. Using a Cramér-Wold device, the CLT follows immediately by the standard argument. ■

Proof of Theorem 3.1:

The result follows from (A.4), (3.8) and Lemma A.3 immediately in view of the continuity of $f_{\mathbf{x}}(\cdot)$. ■

Proof of Theorem 3.2:

In what follows, we label the quantities associated with the bootstrap procedure by the superscript *, which will not be further explained unless misunderstanding may arise.

By design, we have for $\forall \mathbf{i}$

$$\begin{aligned}
\widehat{\boldsymbol{\theta}}_{\mathbf{i}}^* - \widehat{\boldsymbol{\theta}}_{\mathbf{i}} &= \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} y_t^* - \widehat{\boldsymbol{\theta}}_{\mathbf{i}} \\
&= \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{s}(\mathbf{x}_t | \widehat{\boldsymbol{\Theta}}) - \widehat{\boldsymbol{\theta}}_{\mathbf{i}} + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \widehat{\varepsilon}_t \eta_t \\
&= \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \varepsilon_t \eta_t + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} (g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \tilde{\boldsymbol{\Lambda}})) \eta_t \\
&\quad + \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top\right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} (\tilde{s}(\mathbf{x}_t | \tilde{\boldsymbol{\Lambda}}) - \tilde{s}(\mathbf{x}_t | \widehat{\boldsymbol{\Theta}})) \eta_t := \mathbf{A}_{\mathbf{i},1} + \mathbf{A}_{\mathbf{i},2} + \mathbf{A}_{\mathbf{i},3},
\end{aligned}$$

where the second equality follows from the definition of y_t^* .

Note that the term $\mathbf{A}_{\mathbf{i},2}$ has been investigated in the proof of Lemma A.3, and is negligible under the condition $\sqrt{Th}^{p+d/2} \rightarrow 0$. For the term $\mathbf{A}_{\mathbf{i},3}$, we can further write

$$\mathbf{A}_{\mathbf{i},\mathbf{3}} = \left(\sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top \right)^{-1} \sum_{t=1}^T \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^\top \eta_t \cdot (\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} - \hat{\boldsymbol{\theta}}_{\mathbf{i}}) = o_P(\tilde{\boldsymbol{\lambda}}_{\mathbf{i}} - \hat{\boldsymbol{\theta}}_{\mathbf{i}}),$$

where the second equality follows from the fact that $\{\eta_t\}$ are i.i.d. draws from $N(0,1)$ and are independent of the sample.

Therefore, we only need to pay attention to $\mathbf{A}_{\mathbf{i},1}$ below. It suffices to consider $\frac{1}{\sqrt{T}} \sum_{t=1}^T \boldsymbol{\xi}_t^*$, where $\boldsymbol{\xi}_t^* = \frac{1}{\sqrt{h^d}} \tilde{\mathbf{x}}_{\mathbf{i},t} \varepsilon_t \eta_t$. As $\{\eta_t\}$ are i.i.d. draws from $N(0,1)$, it is easy to know that

$$\text{Var}^* \left(\frac{1}{\sqrt{T}} \sum_{t=1}^T \boldsymbol{\xi}_t^* \right) \simeq \text{Var} \left(\frac{1}{\sqrt{T}} \sum_{t=1}^T \frac{1}{\sqrt{h^d}} \tilde{\mathbf{x}}_{\mathbf{i},t} \varepsilon_t \right)$$

in view of the proof of Lemma A.3.

Below, we consider $E^*[\|\boldsymbol{\xi}_1^*\|^2 \cdot I(\|\boldsymbol{\xi}_1^*\| \geq \epsilon\sqrt{T})]$. Write

$$\begin{aligned} E^*[\|\boldsymbol{\xi}_1^*\|^2 \cdot I(\|\boldsymbol{\xi}_1^*\| \geq \epsilon\sqrt{T})] &\leq \{E^*\|\boldsymbol{\xi}_1^*\|^{2+\nu}\}^{\frac{2}{2+\nu}} \left\{E^*[I(\|\boldsymbol{\xi}_1^*\| \geq \epsilon\sqrt{T})]\right\}^{\frac{\nu}{2+\nu}} \\ &\leq \{E^*\|\boldsymbol{\xi}_1^*\|^{2+\nu}\}^{\frac{2}{2+\nu}} \left\{\frac{E^*\|\boldsymbol{\xi}_1^*\|^{2+\nu}}{\epsilon^{2+\nu} T^{\frac{2+\nu}{2}}}\right\}^{\frac{\nu}{2+\nu}} = \frac{1}{\epsilon^\nu T^{\frac{\nu}{2}}} \{E^*\|\boldsymbol{\xi}_1^*\|^{2+\nu}\}^{\frac{1}{2+\nu} \cdot (2+\nu)} \\ &\leq O(1) \frac{1}{\epsilon^\nu T^{\frac{\nu}{2}}} \|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_{\mathbf{i}})_{\varepsilon_1}\|^{2+\nu} \cdot E|\eta_1|^{2+\nu} \cdot \frac{1}{h^{\frac{d}{2} \cdot (2+\nu)}} \\ &= O(1) \frac{1}{\epsilon^\nu (Th^d)^{\frac{\nu}{2}}} \cdot \frac{1}{h^d} \|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_{\mathbf{i}})_{\varepsilon_1}\|^{2+\nu} = O_P(1) \frac{1}{(Th^d)^{\frac{\nu}{2}}}, \end{aligned}$$

where the first inequality follows from Hölder inequality, the second inequality follows from Chebyshev's inequality, the third inequality follows from the definition of E^* and Minkowski inequality, and the last step follows from $\frac{1}{h^d} E\|\tilde{\mathbf{m}}(\mathbf{x}_1 | \tilde{\mathbf{x}}_{\mathbf{i}})_{\varepsilon_1}\|^{2+\nu} = O(1)$ by the proof of Lemma A.3. Consequently,

$$\frac{1}{T} \sum_{t=1}^T E^*[\|\boldsymbol{\xi}_t^*\|^2 \cdot I(\|\boldsymbol{\xi}_t^*\| \geq \epsilon\sqrt{T})] = O_P(1) \frac{1}{(Th^d)^{\frac{\nu}{2}}} = o_P(1),$$

where the last step follows from $Th^d \rightarrow 0$. Therefore, the Lindberg condition is justified. Then the result follows. \blacksquare

Proof of Corollary A.1:

The proof is a simpler version of that presented for Lemma A.3 and Theorem 3.1, therefore it is omitted. \blacksquare

A.5.2 Model (1.2)

Proof of Lemma 3.4:

(1). First, note that provided $0 < x, x_0 < 1$, we have the following two expressions by the following Taylor expansions:

$$\log x = \log x_0 + (x - x_0) \frac{1}{x_0} - (x - x_0)^2 \frac{1}{2(x^*)^2}, \quad (\text{A.22})$$

$$\log(1 - x) = \log(1 - x_0) - (x - x_0) \frac{1}{1 - x_0} - (x - x_0)^2 \frac{1}{2(1 - x^\dagger)^2}, \quad (\text{A.23})$$

where both x^* and x^\dagger lie between x and x_0 .

We are now ready to start our investigation. By (A.22) and (A.23), write

$$\begin{aligned}
& \frac{1}{T} \log L(g(\cdot)) - \frac{1}{T} \log L(\tilde{s}(\cdot | \Theta)) \\
&= -\frac{1}{T} \sum_{t=1}^T (1 - y_t) \{ \log[1 - \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta))] - \log[1 - \Phi_\varepsilon(g(\mathbf{x}_t))] \} \\
&\quad - \frac{1}{T} \sum_{t=1}^T y_t \{ \log \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \log \Phi_\varepsilon(g(\mathbf{x}_t)) \} \\
&= \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \frac{1 - y_t}{1 - \Phi_\varepsilon(g(\mathbf{x}_t))} \\
&\quad + \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \frac{1 - y_t}{2(1 - \Phi_t^\dagger)^2} \\
&\quad - \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \frac{y_t}{\Phi_\varepsilon(g(\mathbf{x}_t))} \\
&\quad + \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \frac{y_t}{2(\Phi_t^*)^2} \\
&= \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \left[\frac{1 - y_t}{1 - \Phi_\varepsilon(g(\mathbf{x}_t))} - \frac{y_t}{\Phi_\varepsilon(g(\mathbf{x}_t))} \right] \\
&\quad + \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \left[\frac{1 - y_t}{2(1 - \Phi_t^\dagger)^2} + \frac{y_t}{2(\Phi_t^*)^2} \right] \\
&:= \mathbb{L}_{T,1} + \mathbb{L}_{T,2}, \tag{A.24}
\end{aligned}$$

where both Φ_t^* and Φ_t^\dagger lie between $\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta))$ and $\Phi_\varepsilon(g(\mathbf{x}_t))$, and the definitions of $\mathbb{L}_{T,1}$ and $\mathbb{L}_{T,2}$ are obvious.

We then consider $\mathbb{L}_{T,1}$ and $\mathbb{L}_{T,2}$ respectively, and start with $\mathbb{L}_{T,1}$. For notational simplicity, we let

$$\begin{aligned}
e_t &= \frac{1 - y_t}{1 - \Phi_\varepsilon(g(\mathbf{x}_t))} - \frac{y_t}{\Phi_\varepsilon(g(\mathbf{x}_t))}, \\
\Delta \Phi_\varepsilon(g(\mathbf{x}_t)) &= \Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta)) - \Phi_\varepsilon(g(\mathbf{x}_t)).
\end{aligned}$$

Simple algebra shows that

$$\begin{aligned}
E[e_t | \mathbf{x}_t] &= 0, \\
E[e_t^2 | \mathbf{x}_t] &= \frac{1}{\Phi_\varepsilon(g(\mathbf{x}_t))[1 - \Phi_\varepsilon(g(\mathbf{x}_t))]}, \\
|\Delta \Phi_\varepsilon(g(\mathbf{x}_t))| &\leq 2. \tag{A.25}
\end{aligned}$$

For any given $\Delta \Phi_\varepsilon(g(\cdot))$, we then consider

$$E \left\| \frac{1}{T} \sum_{t=1}^T \Delta \Phi_\varepsilon(g(\mathbf{x}_t)) e_t \right\|^2$$

$$\begin{aligned}
&= \frac{1}{T^2} \sum_{t=1}^T E[\Delta\Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot e_t^2 \\
&\quad + \frac{1}{T^2} \sum_{t=1}^T (1-t/T) E[\Delta\Phi_\varepsilon(g(\mathbf{x}_1))][\Delta\Phi_\varepsilon(g(\mathbf{x}_{t+1}))] \cdot e_1 e_{t+1} \\
&\leq \frac{4}{T^2} \sum_{t=1}^T E \left[\frac{1}{\Phi_\varepsilon(g(\mathbf{x}_t))[1-\Phi_\varepsilon(g(\mathbf{x}_t))]} \right] \\
&\quad + \frac{1}{T^2} \sum_{t=1}^T (1-t/T) \alpha(t)^{\nu/(2+\nu)} \{E[e_1^{2+\nu}]\}^{2/(2+\nu)} = O\left(\frac{1}{T}\right), \tag{A.26}
\end{aligned}$$

where the inequality follows from Assumption 2.1, and Davydov's inequality and (A.25). By Lemmas A1 and A2 of Newey and Powell (2003), we immediately obtain that

$$\sup_{\tilde{s} \in \mathcal{S}} |\mathbb{L}_{T,1}| = o_P(1). \tag{A.27}$$

We next investigate $\mathbb{L}_{T,2}$. Write

$$\begin{aligned}
\mathbb{L}_{T,2} &= \frac{1}{T} \sum_{t=1}^T [\Delta\Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \left[\frac{1-y_t}{2(1-\Phi_t^\dagger)^2} + \frac{y_t}{2(\Phi_t^*)^2} \right] \\
&\geq \frac{1}{T} \sum_{t=1}^T [\Delta\Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \left\{ \frac{1-y_t}{4[1+(\Phi_t^\dagger)^2]} + \frac{y_t}{2(\Phi_t^*)^2} \right\} \\
&\geq \frac{1}{T} \sum_{t=1}^T [\Delta\Phi_\varepsilon(g(\mathbf{x}_t))]^2 \cdot \left\{ \frac{1-y_t}{4 \cdot 2} + \frac{y_t}{2} \right\} \\
&\geq \frac{1}{8} \cdot \frac{1}{T} \sum_{t=1}^T [\Delta\Phi_\varepsilon(g(\mathbf{x}_t))]^2,
\end{aligned}$$

where the first inequality follows from $\frac{1}{(a+b)^2} \geq \frac{1}{2a^2+2b^2}$ because of $(a+b)^2 \leq 2a^2+2b^2$, the second inequality follows from the fact that Φ_t^* and Φ_t^\dagger lie between $\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \Theta))$ and $\Phi_\varepsilon(g(\mathbf{x}_t))$, and the third inequality follows from that $\frac{1-y_t}{4 \cdot 2} + \frac{y_t}{2} \geq \frac{1}{8}$ because of y_t taking the value of 1 or 0 only.

By the fact that $0 \geq \frac{1}{T} \log L(g(\cdot)) - \frac{1}{T} \log L(\tilde{s}(\cdot | \hat{\Theta}))$, and (A.24) and (A.27), we now conclude that

$$o_P(1) = \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(\tilde{s}(\mathbf{x}_t | \hat{\Theta})) - \Phi_\varepsilon(g(\mathbf{x}_t))]^2 \asymp \frac{1}{T} \sum_{t=1}^T (g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \hat{\Theta}))^2,$$

which completes the proof of this lemma. ■

By Lemma 3.4, it is obvious that

$$\begin{aligned}
&\frac{1}{T} \sum_{t=1}^T (\tilde{s}(\mathbf{x}_t | \tilde{\Lambda}) - \tilde{s}(\mathbf{x}_t | \hat{\Theta}))^2 \\
&= \frac{1}{T} \sum_{t=1}^T (\tilde{s}(\mathbf{x}_t | \tilde{\Lambda}) - g(\mathbf{x}_t) + g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \hat{\Theta}))^2
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2}{T} \sum_{t=1}^T (\tilde{s}(\mathbf{x}_t | \tilde{\mathbf{\Lambda}}) - g(\mathbf{x}_t))^2 + \frac{2}{T} \sum_{t=1}^T (g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \hat{\Theta}))^2 \\
&= O_P(h^{2p}) + o_P(1) = o_P(1),
\end{aligned}$$

where the third equality follows from the first result of Lemma 3.4 and Lemma 3.3.

Before proceeding further, we calculate the partial derivatives of $\log L(\tilde{s}(\cdot | \Theta))$ with respect to each $\theta_{\mathbf{i}}$.

$$\begin{aligned}
\frac{\partial \log L(\tilde{s}(\cdot | \Theta))}{\partial \theta_{\mathbf{i}}} &= \sum_{t=1}^T \left\{ -\frac{(1-y_t) \cdot \phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))}{1 - \Phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))} + \frac{y_t \cdot \phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))}{\Phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))} \right\} \tilde{\mathbf{x}}_{\mathbf{i},t} \\
&= \sum_{t=1}^T \frac{[y_t - \Phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))] \cdot \phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))}{\Phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta)) [1 - \Phi_{\varepsilon}(\tilde{s}(\mathbf{x}_t | \Theta))]} \tilde{\mathbf{x}}_{\mathbf{i},t} \\
&= \sum_{t=1}^T \frac{[y_t - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))] \cdot \phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))}{\Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}})) [1 - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))]} \tilde{\mathbf{x}}_{\mathbf{i},t}, \tag{A.28}
\end{aligned}$$

where the third equality follows from the fact that \mathbf{x}_t can not simultaneous belong to $C_{\mathbf{x}_{0\mathbf{i}},h}$ and $C_{\mathbf{x}_{0\mathbf{j}},h}$ for $\mathbf{i} \neq \mathbf{j}$ by the construction of $\tilde{s}(\cdot | \Theta)$, and $\tilde{\mathbf{x}}_{\mathbf{i},t}$ is the same as that defined in Section 3.2.

Based on $\frac{\partial \log L(\tilde{s}(\cdot | \Theta))}{\partial \theta_{\mathbf{i}}}$ and some tedious calculation, the second order derivative is

$$\begin{aligned}
&\frac{\partial^2 \log L(\tilde{s}(\cdot | \Theta))}{\partial \theta_{\mathbf{i}} \partial \theta_{\mathbf{i}}^{\top}} \\
&= - \sum_{t=1}^T \frac{\phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))^2}{[1 - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))] \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))} \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^{\top} \\
&\quad + \sum_{t=1}^T [y_t - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))] \frac{\phi_{\varepsilon}^{(1)}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))}{[1 - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))] \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))} \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^{\top} \\
&\quad - \sum_{t=1}^T [y_t - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))] \frac{\phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))^2 [1 - 2\Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))]}{[1 - \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))]^2 \Phi_{\varepsilon}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \theta_{\mathbf{i}}))^2} \tilde{\mathbf{x}}_{\mathbf{i},t} \tilde{\mathbf{x}}_{\mathbf{i},t}^{\top}. \tag{A.29}
\end{aligned}$$

Proof of Lemma A.4:

(1). Note that by Lemma A.4, we can write

$$\begin{aligned}
o_P(1) &= \frac{1}{T} \sum_{t=1}^T (g(\mathbf{x}_t) - \tilde{s}(\mathbf{x}_t | \hat{\Theta}))^2 = \frac{1}{T} \sum_{t=1}^T \left(\sum_{\mathbf{x}_t \in C_{\mathbf{x}_{0\mathbf{i}},h}} [g(\mathbf{x}_t) - s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\theta}_{\mathbf{i}})] \right)^2 \\
&= \sum_{\mathbf{i} \in [M]^d} \frac{1}{T} \sum_{t=1}^T I_{\mathbf{i},h}(\mathbf{x}_t) [g(\mathbf{x}_t) - s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\theta}_{\mathbf{i}})]^2 \geq 0,
\end{aligned}$$

where the third equality follows from the fact that \mathbf{x}_t can not simultaneous belong to $C_{\mathbf{x}_{0\mathbf{i}},h}$ and $C_{\mathbf{x}_{0\mathbf{j}},h}$ for $\mathbf{i} \neq \mathbf{j}$. Thus, we must have

$$\frac{1}{T} \sum_{t=1}^T I_{\mathbf{i},h}(\mathbf{x}_t) [g(\mathbf{x}_t) - s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\theta}_{\mathbf{i}})]^2 = o_P(1),$$

which completes the proof of the first result.

(2). By (A.29), we denote

$$\frac{\partial^2 \log L(\tilde{s}(\cdot | \Theta))}{\partial \boldsymbol{\theta}_i \partial \boldsymbol{\theta}_i^\top} := -\mathbf{L}_1(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) + \mathbf{L}_2(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) - \mathbf{L}_3(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)),$$

where the definitions of $\mathbf{L}_j(\cdot)$ for $j = 1, 2, 3$ should be obvious.

First, we consider $\mathbf{L}_2(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i))$ and $\mathbf{L}_3(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i))$. For $\mathbf{L}_2(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i))$, we write

$$\begin{aligned} \frac{1}{T} \tilde{\mathbf{L}}_2(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) &= \frac{1}{T} \mathbf{H} \mathbf{D} \mathbf{L}_2(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) \mathbf{D}^\top \mathbf{H} \\ &= \frac{1}{T} \sum_{t=1}^T [y_t - \Phi_\varepsilon(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) \mathbf{D}^\top \mathbf{H} \\ &= \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(g(\mathbf{x}_t)) - \Phi_\varepsilon(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) \mathbf{D}^\top \mathbf{H} \\ &\quad + \frac{1}{T} \sum_{t=1}^T [y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i)) \mathbf{D}^\top \mathbf{H}, \end{aligned} \tag{A.30}$$

where the definition of $\tilde{\mathbf{X}}_{i,t}(\cdot)$ is obvious.

For the first term on the right hand side of (A.30), we write

$$\begin{aligned} &\left\| \frac{1}{T} \sum_{t=1}^T [\Phi_\varepsilon(g(\mathbf{x}_t)) - \Phi_\varepsilon(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)) \mathbf{D}^\top \mathbf{H} \right\| \\ &\leq \left\{ \frac{1}{T} \sum_{t=1}^T I_{i,h}(\mathbf{x}_t) [\Phi_\varepsilon(g(\mathbf{x}_t)) - \Phi_\varepsilon(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i))]^2 \right\}^{1/2} \\ &\quad \cdot \left\{ \frac{1}{T} \sum_{t=1}^T \|\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)) \mathbf{D}^\top \mathbf{H}\|^2 \right\}^{1/2} \\ &\leq O(1) \left\{ \frac{1}{T} \sum_{t=1}^T I_{i,h}(\mathbf{x}_t) [g(\mathbf{x}_t) - s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)]^2 \right\}^{1/2} \\ &\quad \cdot \left\{ \frac{1}{T} \sum_{t=1}^T \|\mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)) \mathbf{D}^\top \mathbf{H}\|^2 \right\}^{1/2} \\ &= o_P(1), \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz inequality, the second inequality follows from Mean-Value Theorem and the fact that $\phi_\varepsilon(\cdot)$ is uniformly bounded, and the last step follows from the first result of this lemma and (A.14).

For the second term on the right hand side of (A.30), by some tedious algebra and the first result of this lemma, it is not hard to see that

$$\begin{aligned} &\frac{1}{T} \sum_{t=1}^T [y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(s(\mathbf{x}_t | \tilde{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i)) \mathbf{D}^\top \mathbf{H} \\ &= \frac{1}{T} \sum_{t=1}^T [y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{i,t}(g(\mathbf{x}_t)) \mathbf{D}^\top \mathbf{H} \cdot (1 + o_P(1)). \end{aligned}$$

Further, using Assumption 2 and Billingsley's inequality following a procedure similar (but simplified) as in (A.26), and in connection with (A.14), we can show that

$$\left\| \frac{1}{T} \sum_{t=1}^T [y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \mathbf{H} \mathbf{D} \tilde{\mathbf{X}}_{\mathbf{i},t}(g(\mathbf{x}_t)) \mathbf{D}^\top \mathbf{H} \right\| = o_P(1).$$

Based on the above development, we are readily to conclude that

$$\frac{1}{T} \|\tilde{\mathbf{L}}_2(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}}))\| = o_P(1).$$

Similar to the analysis of $\tilde{\mathbf{L}}_2(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}}))$, we can also obtain that

$$\frac{1}{T} \|\mathbf{H} \mathbf{D} \mathbf{L}_3(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}})) \mathbf{D}^\top \mathbf{H}\| = o_P(1).$$

Below, we focus on $\mathbf{L}_1(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}}))$, and write

$$\begin{aligned} & \frac{1}{T} \mathbf{H} \mathbf{D} \mathbf{L}_1(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}})) \mathbf{D}^\top \mathbf{H} \\ &= \frac{1}{T} \mathbf{H} \mathbf{D} \mathbf{L}_1(g(\mathbf{x}_t)) \mathbf{D}^\top \mathbf{H} + \frac{1}{T} \mathbf{H} \mathbf{D} [\mathbf{L}_1(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}})) - \mathbf{L}_1(g(\mathbf{x}_t))] \mathbf{D}^\top \mathbf{H} \\ &= \frac{1}{T} \mathbf{H} \mathbf{D} \mathbf{L}_1(g(\mathbf{x}_t)) \mathbf{D}^\top \mathbf{H} + o_P(1) \\ &= E \left[\frac{[\phi_\varepsilon(g(\mathbf{x}_1))]^2}{[1 - \Phi_\varepsilon(g(\mathbf{x}_1))] \Phi_\varepsilon(g(\mathbf{x}_1))} I_{i,h}(\mathbf{x}_1) \mathbf{H} \mathbf{m}(\mathbf{x}_1 | \tilde{\mathbf{x}}_{\mathbf{i}}) \mathbf{m}(\mathbf{x}_1 | \tilde{\mathbf{x}}_{\mathbf{i}})^\top \mathbf{H} \right] + o_P(1) \\ &= \frac{f_{\mathbf{x}}(\tilde{\mathbf{x}}_{\mathbf{i}}) \phi_\varepsilon(g(\tilde{\mathbf{x}}_{\mathbf{i}}))^2}{[1 - \Phi_\varepsilon(g(\tilde{\mathbf{x}}_{\mathbf{i}}))] \Phi_\varepsilon(g(\tilde{\mathbf{x}}_{\mathbf{i}}))} \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x} + o_P(1), \end{aligned}$$

where the second equality follows from similar steps as those for $\tilde{\mathbf{L}}_2(s(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}, \hat{\boldsymbol{\theta}}_{\mathbf{i}}))$, the third equality follows from a proof similar to those for (A.17), and the fourth equality follows from a development similar to (A.16).

Thus, we can now conclude that for each \mathbf{i}

$$\left\| \frac{1}{T} \mathbf{H} \mathbf{D} \frac{\partial \log L(\tilde{s}(\cdot | \hat{\boldsymbol{\Theta}}))}{\partial \boldsymbol{\theta}_{\mathbf{i}} \partial \boldsymbol{\theta}_{\mathbf{i}}^\top} \mathbf{D}^\top \mathbf{H} - \tilde{\boldsymbol{\Sigma}}_{\mathbf{i}} \right\| = o_P(1),$$

where $\tilde{\boldsymbol{\Sigma}}_{\mathbf{i}}$ is defined in the body of this lemma.

The proof of the second result is now completed.

(3). In view of the fact that $\boldsymbol{\theta}_{\mathbf{i}}^*$ lies between $\hat{\boldsymbol{\theta}}_{\mathbf{i}}$ and $\tilde{\boldsymbol{\lambda}}_{\mathbf{i}}$, the result follows immediately by going through the same procedure as the second result of this lemma.

(4). Write

$$\begin{aligned} & \frac{1}{\sqrt{T} h^d} \mathbf{H} \mathbf{D} \frac{\partial \log L(g(\cdot))}{\partial \boldsymbol{\theta}_{\mathbf{i}}} \\ &= \frac{1}{\sqrt{T} h^d} \sum_{t=1}^T \frac{[y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \phi_\varepsilon(g(\mathbf{x}_t))}{\Phi_\varepsilon(g(\mathbf{x}_t)) [1 - \Phi_\varepsilon(g(\mathbf{x}_t))]} \mathbf{H} \mathbf{D} \tilde{\mathbf{x}}_{\mathbf{i},t} \\ &= \frac{1}{\sqrt{T} h^d} \sum_{t=1}^T \frac{[y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \phi_\varepsilon(g(\mathbf{x}_t))}{\Phi_\varepsilon(g(\mathbf{x}_t)) [1 - \Phi_\varepsilon(g(\mathbf{x}_t))]} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_{\mathbf{i}}) \cdot (1 + o_P(1)) \end{aligned}$$

$$= \frac{1}{\sqrt{Th^d}} \sum_{t=1}^T u_t \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \cdot (1 + o_P(1))$$

where $\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)$ is defined in (A.15), the second equality follows from (A.14), and in the third equality we let

$$\frac{[y_t - \Phi_\varepsilon(g(\mathbf{x}_t))] \cdot \phi_\varepsilon(g(\mathbf{x}_t))}{\Phi_\varepsilon(g(\mathbf{x}_t))[1 - \Phi_\varepsilon(g(\mathbf{x}_t))]} := u_t$$

for notational simplicity. Moreover, as g is defined on $[-a, a]$, it is easy to know that $0 < \mathbf{c} \leq \Phi_\varepsilon(g(\mathbf{x}_t)) \leq \mathbf{C} < 1$. Thus, we can further write

$$|u_t| \leq \phi_\varepsilon(g(\mathbf{x}_t)) \left(\frac{1}{\Phi_\varepsilon(g(\mathbf{x}_t))} \vee \frac{1}{1 - \Phi_\varepsilon(g(\mathbf{x}_t))} \right) = O(1).$$

Also, simple algebra shows that

$$\begin{aligned} E[u_t^2 | \mathbf{x}_t] &= \frac{\phi_\varepsilon(g(\mathbf{x}_t))^2}{\Phi_\varepsilon(g(\mathbf{x}_t))^2 [1 - \Phi_\varepsilon(g(\mathbf{x}_t))]^2} E[y_t^2 - 2y_t \Phi_\varepsilon(g(\mathbf{x}_t)) + \Phi_\varepsilon(g(\mathbf{x}_t))^2 | \mathbf{x}_t] \\ &= \frac{\phi_\varepsilon(g(\mathbf{x}_t))^2}{\Phi_\varepsilon(g(\mathbf{x}_t))^2 [1 - \Phi_\varepsilon(g(\mathbf{x}_t))]^2} E[y_t - 2y_t \Phi_\varepsilon(g(\mathbf{x}_t)) + \Phi_\varepsilon(g(\mathbf{x}_t))^2 | \mathbf{x}_t] \\ &= \frac{\phi_\varepsilon(g(\mathbf{x}_t))^2}{\Phi_\varepsilon(g(\mathbf{x}_t))[1 - \Phi_\varepsilon(g(\mathbf{x}_t))]} \end{aligned}$$

We now move on and write

$$\begin{aligned} &E \left[\left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T u_t \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \right) \left(\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T u_t \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \right)^\top \right] \\ &= \frac{1}{Th^d} \sum_{t=1}^T E \left[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)^\top u_t^2 \right] \\ &\quad + \frac{1}{Th^d} \sum_{t \neq s}^T E \left[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_s | \tilde{\mathbf{x}}_i)^\top u_t u_s \right]. \end{aligned}$$

Similar to (A.19), we have

$$\left\| \frac{1}{Th^d} \sum_{t \neq s}^T E \left[\tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_s | \tilde{\mathbf{x}}_i)^\top u_t u_s \right] \right\| = o_P(1).$$

Thus, we can conclude that

$$\begin{aligned} &E \left[\left(\frac{1}{\sqrt{T}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t \right) \left(\frac{1}{\sqrt{T}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t \right)^\top \right] \\ &= \frac{1}{T} \sum_{t=1}^T E \left[\frac{\phi_\varepsilon(g(\mathbf{x}_t))^2}{\Phi_\varepsilon(g(\mathbf{x}_t))[1 - \Phi_\varepsilon(g(\mathbf{x}_t))]} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i)^\top \right] + o(1) \\ &\rightarrow \frac{f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \phi_\varepsilon(g(\tilde{\mathbf{x}}_i))^2}{\Phi_\varepsilon(g(\tilde{\mathbf{x}}_i))[1 - \Phi_\varepsilon(g(\tilde{\mathbf{x}}_i))]} \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}, \end{aligned}$$

where the last step follows from a procedure similar to (A.16).

Below, we further use small-block and large-block to prove the normality. To employ the small-block and large-block arguments, we partition the set $\{1, \dots, T\}$ into $2k_T + 1$ subsets with large blocks of size l_T and small blocks of size s_T and the last remaining set of size $T - k_T(l_T + s_T)$, where l_T and s_T are selected such that

$$s_T \rightarrow \infty, \quad \frac{s_T}{l_T} \rightarrow 0, \quad \frac{l_T^3}{Th^d} \rightarrow 0, \quad \text{and} \quad k_T \equiv \left\lfloor \frac{T}{l_T + s_T} \right\rfloor,$$

where ν is defined in Assumption 2.1

For $j = 1, \dots, k_T$, define

$$\begin{aligned} \boldsymbol{\xi}_j &= \sum_{t=(j-1)(l_T+s_T)+1}^{jl_T+(j-1)s_T} \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t, & \boldsymbol{\eta}_j &= \sum_{t=jl_T+(j-1)s_T+1}^{j(l_T+s_T)} \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t, \\ \boldsymbol{\zeta} &= \sum_{t=k_T(l_T+s_T)+1}^T \frac{1}{\sqrt{h^d}} \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t. \end{aligned}$$

Note that $\alpha(T) = o(1/T)$ and $k_T s_T / T \rightarrow 0$. By direct calculation, we immediately obtain that

$$\frac{1}{T} E \left\| \sum_{j=1}^{k_T} \boldsymbol{\eta}_j \right\|^2 \rightarrow 0 \quad \text{and} \quad \frac{1}{T} E \|\boldsymbol{\zeta}\|^2 \rightarrow 0.$$

Therefore,

$$\frac{1}{\sqrt{Th^d}} \sum_{t=1}^T \tilde{\mathbf{m}}(\mathbf{x}_t | \tilde{\mathbf{x}}_i) u_t = \frac{1}{\sqrt{T}} \sum_{j=1}^{k_T} \boldsymbol{\xi}_j + o_P(1).$$

By Proposition 2.6 of Fan and Yao (2003), we have as $T \rightarrow 0$

$$\begin{aligned} & \left| E \left[\exp \left(\frac{iw}{\sqrt{T}} \sum_{j=1}^{k_T} \boldsymbol{\xi}_j \right) \right] - \prod_{j=1}^{k_T} E \left[\exp \left(\frac{iw \boldsymbol{\xi}_j}{\sqrt{T}} \right) \right] \right| \\ & \leq 16(k_T - 1)\alpha(s_T) \rightarrow 0, \end{aligned}$$

where i is the imaginary unit. Thus, the Feller condition is fulfilled as follows:

$$\frac{1}{T} \sum_{j=1}^{k_T} E[\boldsymbol{\xi}_j \boldsymbol{\xi}_j^\top] \rightarrow \frac{f_{\mathbf{x}}(\tilde{\mathbf{x}}_i) \phi_\varepsilon(g(\tilde{\mathbf{x}}_i))^2}{\Phi_\varepsilon(g(\tilde{\mathbf{x}}_i)) [1 - \Phi_\varepsilon(g(\tilde{\mathbf{x}}_i))]} \int_{[-1,1]^d} \mathbf{m}(\mathbf{x} | \mathbf{0}) \mathbf{m}(\mathbf{x} | \mathbf{0})^\top d\mathbf{x}.$$

Also, we note that

$$\begin{aligned} E[\|\boldsymbol{\xi}_1\|^2 \cdot I(\|\boldsymbol{\xi}_1\| \geq \epsilon\sqrt{T})] &\leq \{E\|\boldsymbol{\xi}_1\|^4\}^{\frac{1}{2}} \left\{ E[I(\|\boldsymbol{\xi}_1\| \geq \epsilon\sqrt{T})] \right\}^{\frac{1}{2}} \\ &\leq \{E\|\boldsymbol{\xi}_1\|^4\}^{\frac{1}{2}} \left\{ \frac{E\|\boldsymbol{\xi}_1\|^4}{\epsilon^4 T^2} \right\}^{\frac{1}{2}} \\ &= \frac{1}{\epsilon^2 T} \{E\|\boldsymbol{\xi}_1\|^4\}^{\frac{1}{4} \cdot 4} = O(1) \frac{l_T^4}{\epsilon^2 T h^d}, \end{aligned}$$

where the first inequality follows from Hölder inequality, the second inequality follows from Chebyshev's inequality, and the last step follows from Minkowski inequality. Consequently,

$$\frac{1}{T} \sum_{j=1}^{k_T} E[\|\boldsymbol{\xi}_j\|^2 \cdot I(\|\boldsymbol{\xi}_j\| \geq \epsilon\sqrt{T})] = O\left(\frac{k_T l_T^4}{T \cdot T h^d}\right) = O\left(\frac{l_T^3}{T h^d}\right) = o(1),$$

which is the Lindberg condition. Using a Cramér-Wold device, the CLT follows immediately by the standard argument. ■

Proof of Corollary 3.3:

By the first order condition, we have

$$0 = \frac{\partial \log L(\tilde{s}(\cdot | \widehat{\boldsymbol{\Theta}}))}{\partial \boldsymbol{\theta}_i} = \frac{\partial \log L(s(\cdot | \tilde{\mathbf{x}}_i, \widehat{\boldsymbol{\theta}}_i))}{\partial \boldsymbol{\theta}_i},$$

where the second equality follows from (A.28).

Using Taylor expansion, we have

$$0 = \frac{\partial \log L(s(\cdot | \tilde{\mathbf{x}}_i, \tilde{\boldsymbol{\lambda}}_i))}{\partial \boldsymbol{\theta}_i} + \frac{\partial^2 \log L(s(\cdot | \tilde{\mathbf{x}}_i, \boldsymbol{\theta}_i^*))}{\partial \boldsymbol{\theta}_i \partial \boldsymbol{\theta}_i^\top} (\widehat{\boldsymbol{\theta}}_i - \tilde{\boldsymbol{\lambda}}_i)$$

where $\boldsymbol{\theta}_i^*$ lies between $\widehat{\boldsymbol{\theta}}_i$ and $\tilde{\boldsymbol{\lambda}}_i$, and the second equality follows from Lemma 3.3 and the continuity of ϕ_ϵ and Φ_ϵ .

Thus, by Lemma A.4, the result follows immediately. ■

Proof of Theorem 3.4:

The proof follows from (A.28) and (A.29) and a procedure very similar to that given in Theorem 3.2. ■