ECONSTOR Make Your Publications Visible.

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Fan, Jianqing; Wang, Weining; Zhao, Yue

Working Paper Conditional nonparametric variable screening by neural factor regression

cemmap working paper, No. CWP17/24

Provided in Cooperation with: Institute for Fiscal Studies (IFS), London

Suggested Citation: Fan, Jianqing; Wang, Weining; Zhao, Yue (2024) : Conditional nonparametric variable screening by neural factor regression, cemmap working paper, No. CWP17/24, Centre for Microdata Methods and Practice (cemmap), London, https://doi.org/10.47004/wp.cem.2024.1724

This Version is available at: https://hdl.handle.net/10419/302890

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Conditional nonparametric variable screening by neural factor regression

Jianqing Fan Weining Wang Yue Zhao

The Institute for Fiscal Studies Department of Economics, UCL

cemmap working paper CWP17/24



Conditional nonparametric variable screening by neural factor regression

Jianqing Fan

Operations Research and Financial Engineering, Princeton University Weining Wang

Department of Economics, Econometrics and Finance, University of Groningen

Yue Zhao

Department of Mathematics, University of York

August 17, 2024

Abstract

High-dimensional covariates often admit linear factor structure. To effectively screen correlated covariates in high-dimension, we propose a conditional variable screening test based on non-parametric regression using neural networks due to their representation power. We ask the question whether individual covariates have additional contributions given the latent factors or more generally a set of variables. Our test statistics are based on the estimated partial derivative of the regression function of the candidate variable for screening and a observable proxy for the latent factors. Hence, our test reveals how much predictors contribute additionally to the non-parametric regression after accounting for the latent factors. Our derivative estimator is the convolution of a deep neural network regression estimator and a smoothing kernel. We demonstrate that when the neural network size diverges with the sample size, unlike estimating the regression function itself, it is necessary to smooth the partial derivative of the neural network estimator to recover the desired convergence rate for the derivative. Moreover, our screening test achieves asymptotic normality under the null after finely centering our test statistics that makes the biases negligible, as well as consistency for local alternatives under mild conditions. We demonstrate the performance of our test in a simulation study and two real world applications.

Keywords: Neural networks, factor model, non-parametric regression, non-parametric tests, functional of derivatives, high-dimensionality.

1 Introduction

1.1 Background

Variable screening is a powerful tool to expeditiously identify the set of predictors that potentially affect the regression outcome (Fan and Lv, 2008). It can reduce a very large number of predictors to a smaller and more manageable set. Then, on this reduced set, one can apply some more refined but computationally demanding variable selection methods such as Lasso, SCAD, Danzig selector (Tibshirani, 1996; Fan and Li, 2001; Candes and Tao, 2007; Fan et al., 2020) and their non-parametric counterpart FAST-NN (Fan and Gu, 2023). Conditional marginal screening in parametric regression (Barut et al., 2016) further augments the screening by conditioning on a *known* set of useful predictors to reduce the impact of the correlations among the predictors, thus making the important predictors more visible and reducing the false positive and false negative rates in the vanilla screening method. Despite the aforementioned advances, conducting variable screening or selection in non-parametric regression with high-dimensional inputs remains a challenge due to curse of dimensionality, and deep learning offers a promising solution thanks to its ability to adapt to unknown low-dimensional structures in multivariate non-parametric regression.

Deep learning has achieved tremendous empirical successes in numerous applications (LeCun et al., 2015; Goodfellow et al., 2016), for instance, in high-dimensional problems such as image recognition (Simonyan and Zisserman, 2015; Krizhevsky et al., 2017), deep reinforcement learning (Mnih et al., 2015), and large language models (Kasneci et al., 2023; Thirunavukarasu et al., 2023). There is now also a growing literature justifying theoretically the benefit of depth in deep neural networks (Telgarsky, 2016; Yarotsky, 2017; Elbrächter et al., 2021) and their power in alleviating the curse of dimensionality in non-parametric regression via algorithmic learning of unknown low-dimensional structures

within complex functions (Schmidt-Hieber, 2020; Kohler and Langer, 2021). Now being a component of the standard toolbox for statisticians, deep neural networks may nevertheless not be efficient if the dimension of the predictors is very high due to the fundamental limit of multivariate non-parametric regression. Recently Fan and Gu (2023) proposed a factor-augmented sparse throughput regression model that simultaneously leveraged the aforementioned adaptivity of deep neural networks and a factor model on the predictors to facilitate variable selection. To complement the variable selection effort in Fan and Gu (2023), Dinh and Ho (2020), and Ho et al. (2023), in this paper, we investigate the issue of conditional variable screening with deep neural networks when facing potentially very high-dimensional inputs.

We will assess conditional contribution of a candidate variable for screening by examining its partial derivative in the multivariate non-parametric regression function with a given set of variables and construct our screening test statistics on the moment generating function (MGF) of the smoothed partial derivative of the regression function estimator. We chose a deep neural network as our regression estimator due to its aforementioned algorithmic adaptation to the low-dimensional structure. Thus, quantifying and improving the performance of derivatives of deep neural networks are integral to our study, and also form an interesting topic of its own right given the importance of derivative estimation in non-parametric regression across diverse research domains and practical applications (Gijbels and Goderniaux, 2005; Rondonotti et al., 2007; Horel and Giesecke, 2020). Traditional non-parametric derivative estimation in general follows one of the following three methods: empirical derivative-type estimation (Müller et al., 1987; De Brabanter et al., 2013; Liu and De Brabanter, 2020), kernel/local polynomial-type estimation (Gasser and Müller, 1984; Fan and Gijbels, 1996), and series/spline-type estimation (Stone, 1985; Zhou and Wolfe, 2000). However, the first method heavily relies on the existence of an order among the predictor samples and hence naturally applies when the predictor dimension is just one, and the last two methods, just like their original regression estimation counterparts, suffer from the curse of dimensionality problem when facing predictors of a moderate dimension.

Despite their deteriorated performances posed by high dimensionality, the traditional derivative estimation methods are relatively amenable to theoretical analysis due to their closed-form solutions. Excluding the empirical derivative method, which avoids explicitly fitting a regression function, the closed-form solutions for the last two methods can be attributed to the close connection between estimating the original regression functions and their accompanying derivatives. For instance, in the kernel and the spline methods, a closed-form derivative estimator can simply be obtained as the derivative of the original regression function estimator (see, for instance, the discussion between Eqs. (2) and (3) in Zhou and Wolfe (2000)). Thus, unsurprisingly, in these methods the quality of the derivative estimation closely follows the quality of the original regression function estimation. However, for deep neural networks, estimating the *derivative* of a regression function can be quite different from the task on the regression function itself due to its smoothness.

Take for example any candidate regression function estimator m within the canonical neural network class, precisely defined in (6) later, built from the popular ReLU activation function. The first order partial derivatives of m are necessarily piecewise constant due to the piecewise linear nature of the ReLU function, and the second order partial derivatives of m are necessarily zero almost everywhere, *irrespective* of the underlying truth that the function m may attempt to recover. Numerically, this is easily demonstrable through standard software packages such as PyTorch. In addition, the first order partial derivatives of m could exhibit convergence behaviors qualitatively different from the convergence behavior of m itself as we will explain in Section 2.3. Such phenomena are in sharp contrast to the traditional non-parametric estimators, and can in part be attributed to the features of deep neural networks: they are highly non-linear and lack easily interpretable closedform expressions. Moreover, existing asymptotic results on functionals acting on deep neural networks almost invariably assume some continuity of the functionals with respect to their inputs. However, when such functionals in effect act on the derivatives of neural networks, the assumed continuities could break down due to the aforementioned different convergence behaviors of the said derivatives. We will explain and address this discrepancy as we progress through the paper.

1.2 Our method and contribution

Our main contribution in this paper is a non-parametric conditional variable screening test using deep neural networks when the ambient dimension of the inputs is potentially very high. In addition to complementing the variable selection methods, our contribution is also an advance over the aforementioned paper by Barut et al. (2016) in that we work with non-parametric regression and our conditioning variables are not known in advance but instead are extracted from the inputs with the help of the factor model as in Fan and Liao (2022) and Fan and Gu (2023). We note that Horel and Giesecke (2020) has also conducted screening test based on partial derivatives (in low dimensions). However their study focused on single-layer neural networks, and hence does not benefit from the power nor reveal the intricacies of deep neural networks in derivative estimation.

In addition to the general procedures outlined above, we also make the following contributions which could be of independent interest:

1) We rigorously derive the size and power of our proposed test statistics employing non-linear functionals of truly deep neural networks whose architecture can become more complex as the sample size increases. For computational ease, we propose a simple variance estimator whose properties we properly characterize. Our resulting test statistics are straightforward to compute and sufficiently precise to accommodate local alternatives. Moreover, we address the aforementioned continuity issue of functionals acting on the derivatives of neural networks by proposing improved estimation of the said derivatives which naturally leads to our next contribution.

2) We exploit the regression function algorithmically learned by deep neural networks in order to recover high-quality derivative estimation. To achieve this goal, we employ a smoothing technique on deep neural network estimators to regularize their derivatives. Moreover, out method is applicable to regularize the derivatives of alternative machine learning techniques, thus paving the way for their application in derivative estimation. Last but not least, although we focus on (conditional) marginal screening in the present paper, our method can be generalized easily to test higher-order effect and variable interaction using higher-order and mixed derivatives respectively.

Through simulation studies, we illustrate the favorable size and power performance of our test statistics, and the benefit of the smoothing operation in generating accurate derivative estimators. We also highlight the potential of our test statistics as a viable tool for model specification in nonlinear factor models via two empirical applications.

1.3 Notations, conventions, and manuscript organization

Let \mathbb{N} denote the set of positive integers. For a vector $\mathbf{v} = (v_1, \ldots, v_d)^\top \in \mathbb{R}^d$, we let $\|\mathbf{v}\|_p = (\sum_{k=1}^d |v_i|^p)^{1/p}$ be the ℓ_p norm of \mathbf{v} . For a matrix $\mathbf{A} = (a_{j,k})_{1 \leq j \leq l, 1 \leq k \leq m} \in \mathbb{R}^{l \times m}$, we define the operator norm $\|\mathbf{A}\|_{\text{op}} = \max_{\mathbf{v} \in \mathbb{R}^m: \|\mathbf{v}\|_2 = 1} \|\mathbf{A}\mathbf{v}\|_2$. We let C denote an (absolute) constant that may change for each occurrence, and let " \leq " denote an inequality that holds up to such a multiplicative factor C; moreover, let c, C and M with super/subscripts denote constants with particular (though often non-specified) values. Limits are taken as $n \to \infty$ unless otherwise stated. For positive number sequences $(a_n : n \geq 1)$ and $(b_n : n \geq 1)$, we denote $a_n \leq b_n$ if there exists a positive constant C such that $a_n/b_n \leq C$ (for all n), and denote $a_n = \mathcal{O}(b_n)$ (resp. $a_n \sim b_n$) if $a_n/b_n \to 0$ (resp. $a_n \leq b_n$ and $b_n \leq a_n$). We use " \rightarrow_d " to denote convergence in distribution. Finally, let $\|\cdot\|_{L_2} = \|\cdot\|_{L_2(\mathbb{P})}$ denote the L_2 norm of the argument function with respect to the measure \mathbb{P} to be formally introduced in Section 2.1, so $\|f\|_{L_2} = \{\int f^2 d\mathbb{P}\}^{1/2}$. Sections in the supplement are labelled alphabetically.

We organize the remainder of our paper as follows. Section 2 specifies our regression model and screening test, and develop our derivative estimators and test statistics. Section 3 depicts the accompanying theoretical properties of the derivative estimators and test statistics from Section 2. Section 4 presents a simulation study. Section 5 applies our test to an empirical example in asset pricing and another in macroeconomics time series. Section 6 concludes and suggests several extensions. Additional results for the empirical application, proofs and supporting details are deferred to the supplementary materials.

2 Derivative estimator and test statistics

2.1 A conditional screening test through latent factors

Our screening test aims to tackle the high-dimensional regime where the ambient dimension d of our observed predictors $X \in \mathbb{R}^d$ can increase with the sample size n. Even with the remarkable capacity of deep neural networks to represent complex functions, our task is still infeasible if d grows too fast. To address this issue, we assume that X admits the following factor model (Fan and Gu, 2023):

$$\boldsymbol{X} = \boldsymbol{B}\boldsymbol{F} + \boldsymbol{u}, \quad \mathbb{E}(\boldsymbol{u}|\boldsymbol{F}) = \boldsymbol{0}, \tag{1}$$

where $\mathbf{F} \in \mathbb{R}^r$ is a vector of latent factors, $\mathbf{B} \in \mathbb{R}^{d \times r}$ (usually $d \gg r$ for the high-dimensional \mathbf{X}) is the factor loading matrix, and \mathbf{u} is the vector of the idiosyncratic noises. We let $(Y, \mathbf{F}, \mathbf{X})$ have joint distribution \mathbb{P} which also determines the distribution of \mathbf{u} by (1).

The conditional marginal screening is to see whether a component X_j of \mathbf{X} has additional contributions to the response variable Y given \mathbf{F} . Assume that $(Y, \mathbf{F}, \mathbf{X})$ is a $\mathbb{R} \times [-b, b]^{r+d}$ valued random vector. Then, the problem involves the working regression function $m_0(\cdot)$: $[-b, b]^{r+1} \to \mathbb{R}$ defined through

$$Y = m_0(\mathbf{F}, X_j) + \epsilon, \quad \text{where } m_0(\mathbf{F}, X_j) \equiv \mathbb{E}[Y|\mathbf{F}, X_j].$$
⁽²⁾

The additional contribution of variable X_j is measured through the partial derivative

$$m_{0,j} = m_{0,j}(\boldsymbol{f}, x_j) \equiv \frac{\partial}{\partial x_j} m_0(\boldsymbol{f}, x_j).$$

We impose a blanket notational convention that a function with subscript j denotes the partial derivative with respect to the last argument, which will almost always be x_j , while keeping the other argument (in this case, f) fixed. Then, to examine whether X_j has additional contribution, we propose the null and alternative hypotheses

$$H_0: \text{for all } \boldsymbol{f}, x_j \in [-b, b]^{r+1}, \ m_{0,j}(\boldsymbol{f}, x_j) = 0, \text{ against}$$
(3)

$$H_A$$
: for some $\boldsymbol{f}, x_j \in [-b, b]^{r+1}$, we have $m_{0,j}(\boldsymbol{f}, x_j) \neq 0.$ (4)

The function m_0 is akin to $\mathbb{E}[Y|X_j]$ in condition C in Fan et al. (2011) on unconditional non-parametric screening, but here we exercise finer control over the potential role of X_j by introducing f in $m_0 = m_0(f, x_j)$. Note that working with m_0 involves estimating the latent factors F, which will be discussed in Section 2.2. By Remark A.1 in the supplement, the screening hypotheses for X_j in (3) and (4) are also equivalent to a significance test for the *j*-th idiosyncratic term u_j . We deliberately avoid working with the full conditional expectation of Y, namely $m_0^*(F, X) \equiv \mathbb{E}[Y|F, X]$, because this is likely infeasible if the dimension of X is too high without additional structure. In contrast, the total number of variables r + 1 in our working regression function m_0 and its derivative $m_{0,j}$ is much less than d, which allows us to circumvent the problem of testing a potentially very large number of coordinates of X simultaneously (Fan and Gu, 2023). Nevertheless, our model retains its validity even when the regression outcome Y involves more coordinates of X, as we will explain in Remark A.2.

Although we shall focus on the high-dimensional X case with growing d, our screening test also easily accommodates the low-dimensional X case, as we will comment in Section D.1. At the other end of the spectrum, our findings can be generalized to testing X over a fixed or expanding set of coordinates that form a subset of $\{1, \dots, d\}$, as we will briefly outline in Section 6.

2.2 Initial regression estimator through diversified projection

We focus on multi-layer feed-forward neural networks that are fully connected between adjacent layers (Anthony and Bartlett, 1999, p. 75). In this paper we exclusively consider the Rectified Linear Unit (ReLU) activation function defined as $\sigma = \sigma_{\text{ReLU}}(x) = \max\{x, 0\}$ due to its widespread popularity (Goodfellow et al., 2016, Section 6.1).

For now, we consider neural network functions with a generic input dimension \overline{d} . We finalize the structure of our neural networks by specifying the tuple (L, \mathbf{k}) where $L \in \mathbb{N}$ represents the number of hidden layers and the width vector $\mathbf{k} = (\overline{d}, k_1, \ldots, k_L, k_{L+1} = 1) \in$ \mathbb{N}^{L+2} specifies the number of nodes (i.e., neurons) in each hidden layer and the input/output dimensions. More precisely, such a neural network function $f : \mathbb{R}^{\overline{d}} \to \mathbb{R}$ with architecture (L, \mathbf{k}) is given by

$$f(\boldsymbol{x}) = \mathcal{L}_{L+1} \circ \overline{\sigma}_L \circ \mathcal{L}_L \circ \overline{\sigma}_{L-1} \circ \cdots \circ \mathcal{L}_2 \circ \overline{\sigma}_1 \circ \mathcal{L}_1(\boldsymbol{x});$$
(5)

here $L_{\ell}(\boldsymbol{z}) = \boldsymbol{w}_{\ell}\boldsymbol{z} + \boldsymbol{b}_{\ell}$ is an affine map with weight matrix $\boldsymbol{w}_{\ell} \in \mathbb{R}^{k_{\ell} \times k_{\ell-1}}$ if $\ell \ge 2$, $\boldsymbol{w}_{1} \in \mathbb{R}^{k_{1} \times \overline{d}}$ and bias vector $\boldsymbol{b}_{\ell} \in \mathbb{R}^{k_{\ell}}$, and the function $\overline{\sigma}_{\ell} : \mathbb{R}^{k_{\ell}} \to \mathbb{R}^{k_{\ell}}$ applies the ReLU activation function σ entry-wise. Additionally, we will truncate the (univarite) output of f at a prespecified constant level M > 0 with the truncation operator $T_{M}(x) \equiv \operatorname{sgn}(x) \min\{|x|, M\}$. We denote such a collection of truncated f by $\mathcal{F}_n(\overline{d})$ where for brevity of notation we only retain the dependence on the input dimension \overline{d} :

$$\mathcal{F}_n(\overline{d}) = \{T_M(f) : f \text{ is of the form (5) with } L \text{ hidden layers and width vector } \mathbf{k}\}, \quad (6)$$

where both L and \mathbf{k} can scale with the sample size n.

Next, we address the latency issue of the unobserved factor F. For a given diversified weight $\boldsymbol{w} \in \mathbb{R}^d$, Fan and Liao (2022) note from (1) that within the projection $\boldsymbol{w}^T \boldsymbol{X}$, the projected idiosyncratic terms $\boldsymbol{w}^T \boldsymbol{u}$ will be negligible in high-dimension due to the law of the averages, so $\boldsymbol{w}^T \boldsymbol{X}$ yields an approximate linear combination of F. To yield r factors, we need at least r projections. Since r is unknown, one specifies an upper bound \bar{r} of r.

Let $\mathbf{W} \in \mathbb{R}^{d \times \overline{r}}$ be a pre-trained *diversified projection matrix* as termed by Definition 3 in Fan and Gu (2023), and let $\mathbf{H} \equiv d^{-1} \mathbf{W}^{\top} \mathbf{B} \in \mathbb{R}^{\overline{r} \times r}$. Then, by (1), we have

$$\widetilde{\boldsymbol{F}} \equiv d^{-1} \mathbf{W}^{\top} \boldsymbol{X} = \mathbf{H} \boldsymbol{F} + d^{-1} \mathbf{W}^{\top} \boldsymbol{u} \approx \mathbf{H} \boldsymbol{F} \in \mathbb{R}^{\overline{r}},$$

by the law of the averages. Hence, $\mathbf{H}^{\dagger} \widetilde{F} \approx F$ under some appropriate conditions, where \mathbf{H}^{\dagger} is the pseudo-inverse of \mathbf{H} . We call \widetilde{F} as the *diversified factor*, which is observable and a proxy of the latent factor F. In practice, acquiring \mathbf{W} in advance is necessary, either through domain expertise or data-driven methods. For instance, Proposition 1 in Fan and Gu (2023) proposes estimating \mathbf{W} via pretraining, where a tiny portion of size approximately $\log(n)$ of the samples is reserved to extract the top \overline{r} principal components in order to construct a \mathbf{W} that will satisfy Assumption B.2 with high probability. Henceforth we shall assume that \mathbf{W} is pre-determined, exogenous and fixed.

With the proxies of the latent factors, for each given j, the coordinate of the candidate variable for screening, we can compute $\{Y_i, \tilde{F}_i, X_{i,j}\}_{i=1}^n$ based on the observed data, where $\tilde{F}_i = d^{-1} \mathbf{W}^\top \mathbf{X}_i$. Then we fit the neural network regression

$$\widehat{g}_n = \operatorname{argmin}_{g \in \mathcal{F}_n(\overline{r}+1)} \frac{1}{n} \sum_{i=1}^n \{Y_i - g(\widetilde{F}_i, X_{i,j})\}^2 = \operatorname{argmin}_{g \in \mathcal{F}_n(\overline{r}+1)} \mathbb{P}_n \ell(\cdot; g)$$
(7)

where $\ell(y, \tilde{f}, x_j; g) = \{y - g(\tilde{f}, x_j)\}^2$ is the square loss and \mathbb{P}_n is the empirical distribution. Intuitively, \hat{g}_n is a neural network estimation of

$$g_0 = g_0(\widetilde{\boldsymbol{f}}, x_j) \equiv m_0(\mathbf{H}^{\dagger} \boldsymbol{f}, x_j).$$
(8)

Indeed, Lemma B.2 shows that g_0 approximates m_0 well.

2.3 Initial test statistic and its improvements

To test our screening hypotheses (3) and (4), we start from $\hat{g}_{n,j}$, which by the convention in Section 2.1 is the partial derivative of \hat{g}_n in (7) with respect to x_j : $\hat{g}_{n,j}(\tilde{f}, x_j) = \partial \hat{g}_n(\tilde{f}, x_j)/\partial x_j$. Now, $\hat{g}_{n,j}$ can be regarded as an estimator of the derivative $m_{0,j}$. We then consider the following *initial* MGF/exponentially tilted test statistic:

$$\eta_t(\widehat{g}_n) = \mathbb{P}_n\{\exp(t\,\widehat{g}_{n,j}) - 1\}.$$
(9)

Partly owing to the fact that a MGF uniquely characterizes the distribution of the underlying random variable (specifically when considering the MGF over an interval including zero), MGF-based tests are popular in the literature (Epps et al., 1982; Baringhaus et al., 2017). For instance, they have been extensively employed in testing (multivariate) normality (Ebner and Henze, 2020). Under the null hypothesis H_0 : $m_{0,j} = 0$, we expect $\hat{g}_{n,j}$ to be close to zero and so $\eta_t(\hat{g}_n)$ is also centered around zero.

To enhance the quality of both the derivative estimator \hat{g}_n and the test statistic $\eta_t(\hat{g}_n)$, we will further conduct the following refinements sequentially:

(a) **Smoothing**. As already alluded to in the introduction, in derivative estimation, the straightforward plugin estimator $\hat{g}_{n,j}$ derived from \hat{g}_n in (7) may perform poorly due to irregularities of the neural network functions. For example, for the *L*-time iterated sawtooth function ζ_L (see Lemma 2.4 in Telgarsky (2015) and Figure E.1) which is implementable by a neural network of no more than *L* hidden layers and three nodes

per layer, we have $\|\dot{\zeta}_L\|_{L_2} \ge C2^L \|\zeta_L\|_{L_2}$, which implies that the quality of derivative estimation can become significantly worse than that of regression function estimation as the depth L increases, a regime precisely of interest for deep neural networks. This phenomenon results from the increasingly oscillatory behaviour of the neural network functions as the depth increases, and hence is not tied specifically to the ReLU activation function. To address this issue, we will refine the initial estimator $\hat{g}_{n,j}$ to obtain a smoothed derivative estimator $\hat{g}_{n,j}^s$, which will allow us to recover a faster convergence rate; see Theorem 3.1 and Theorem 3.4; this will in turn address the continuity issue arising from second-order terms in our MGF test statistics acting on the derivatives of neural networks. Throughout this paper, we adhere to the convention that the superscript "s" in upright font denotes the smoothed variant of the preceding function. We will provide more detailed calculation for our observation on the sawtooth function and more comprehensive rationale behind the smoothing operation in Section A.4.

- (b) Centering. We must further refine the initial \$\tilde{g}_{n,j}^{s}\$ to arrive at \$\tilde{g}_{n,j}^{s}\$ in (14) in order to center our test statistics. This debiasing step is analogous to the concept of targeted machine learning in the literature (van der Laan et al., 2011; van der Laan and Rose, 2018), and also relates to other research on debiased machine learning (Quintas-Martinez et al., 2022; Kennedy, 2022). However, we will justify this refinement step independently without directly referencing the targeted/debiased machine learning literature.
- (c) **Truncation**. To treat the technical possibility of the unboundedness of $\check{g}_{n,j}^{s}$, we will truncate $\check{g}_{n,j}^{s}$ appropriately in our final test statistics.
- (d) **Uniformity over** t. We will extend the fixed-t statistic to the statistics aggregated over a range of t to enhance the power of the test (Fan et al., 2001, 2015).

We focus solely on the smoothing step (a) in this section, and will introduce step (b) in Section 2.5, and both (c) and (d) in Section 2.6.

We first describe the smoothing operation applied to a generic function $g : \mathbb{R}^{\bar{r}+1} \to \mathbb{R}$ (or analogously $m : \mathbb{R}^{r+1} \to \mathbb{R}$). Let K be a univariate continuously differentiable kernel supported on [-1, 1] with derivative \dot{K} . Then, for the generic function g, denote its smoothed version in the variable x_j by $g^s = g^s(\tilde{f}, x_j) = \int g(\tilde{f}, z) K_h(x_j - z) dz$, where $K_h(\cdot) = K(\cdot/h)/h$ with a bandwidth parameter h. Now, the smoothed function g^s becomes differentiable in x_j everywhere with the partial derivative given by

$$g_j^{\rm s}(\widetilde{\boldsymbol{f}}, x_j) = \frac{\partial}{\partial x_j} g^{\rm s}(\widetilde{\boldsymbol{f}}, x_j) = \int_{x_j-h}^{x_j+h} g(\widetilde{\boldsymbol{f}}, z) \dot{K}_h(x_j-z) \mathrm{d}z = \frac{1}{h} \int_{-1}^{1} g(\widetilde{\boldsymbol{f}}, x_j-ah) \dot{K}(a) \mathrm{d}a.$$
(10)

Accordingly, we let our initial smoothed derivative estimator based on \hat{g}_n be $\hat{g}_{n,j}^s$. In practice, we can select the bandwidth h through cross-validation; see our Remark A.4. See Figure E.2 for the performance of the smoothed derivative estimator across varying bandwidths in a simple, low-dimensional example detailed in Section D.2.

2.4 Estimating the score function

To test the screening hypotheses (3) and (4), we proceed to estimate the score function corresponding to the statistic $\eta_t^{s}(\hat{g}_n) \equiv \mathbb{P}_n\{\exp(t\,\hat{g}_{n,j}^s) - 1\}$ now refined over (9). This score function $\alpha_{t,n}^* : \mathbb{R}^{\bar{r}+1} \to \mathbb{R}$ satisfies, under the null and for any $L_2(\mathbb{P})$ -integrable α , the equality $\int \alpha \alpha_{t,n}^* d\mathbb{P} = \int_{\Omega_h} \alpha_j^s(\tilde{f}, x_j) d\mathbb{P}$, where α_j^s is the smoothed derivative of α and $\Omega_h = \{\omega \in \Omega : X_j(\omega) \in \mathcal{B}_h\}$ is the interior sample space corresponding to the interior set $\mathcal{B}_h = [-b + h, b - h]$. (Such an $\alpha_{t,n}^*$ exists because it is the Riesz representer of the directional derivative functional of our test statistic, as we will formally explain in Section 3.2.1.) Therefore, the mean-square error (MSE) can be expressed as

$$\mathbb{E}(\alpha - \alpha_{t,n}^*)^2(\widetilde{\boldsymbol{F}}, X_j) = \mathbb{E}\alpha^2(\widetilde{\boldsymbol{F}}, X_j) - 2\mathbb{E}(\alpha \alpha_{t,n}^*)(\widetilde{\boldsymbol{F}}, X_j) + \mathbb{E}\alpha_{t,n}^{*2}(\widetilde{\boldsymbol{F}}, X_j)$$

^{under} ${}^{H_0}\mathbb{E}\alpha^2(\widetilde{\boldsymbol{F}}, X_j) - 2\int_{\Omega_h} \alpha_j^s(\widetilde{\boldsymbol{f}}, x_j) d\mathbb{P} + \mathbb{E}\alpha_{t,n}^{*2}(\widetilde{\boldsymbol{F}}, X_j).$ (11)

Discarding the last term $\mathbb{E}[\alpha_{t,n}^{*2}(\widetilde{\mathbf{F}}, X_j)]$, which does not depend on α , the MSE can be estimated by the empirical loss

$$\widehat{R}_{n}^{\text{null}}(\alpha) = \frac{1}{n} \sum_{i=1}^{n} \alpha^{2}(\widetilde{F}_{i}, X_{i,j}) - 2\frac{1}{n} \sum_{i \in \mathcal{I}_{h}} \alpha_{j}^{\text{s}}(\widetilde{F}_{i}, X_{i,j}),$$
(12)

where $\mathcal{I}_h = \{i \in \{1, \ldots, n\} : X_{i,j} \in \mathcal{B}_h\}$ is the interior index set. Hence, we estimate $\alpha_{t,n}^*$ by

$$\widehat{\alpha}_n = \operatorname{argmin}_{\alpha \in \mathcal{F}_n(\overline{r}+1)} \widehat{R}_n^{\operatorname{null}}(\alpha), \tag{13}$$

which is a neural network estimator of the score function.

We tailor the loss function \hat{R}_n^{null} and the estimator $\hat{\alpha}_n$ to the null hypothesis (3). One notable computational advantage is that \hat{R}_n^{null} and hence $\hat{\alpha}_n$ are independent of both the response and the value of t: A single optimization suffices to yield $\hat{\alpha}_n$ for all t. However, the potential trade-off is a bias of $\hat{\alpha}_n$ with respect to $\alpha_{t,n}^*$ induced under the alternative hypothesis, which will be characterized in Theorem 3.2. Numerically, computing the righthand side in (12) is straightforward, as detailed in Lemma A.2 in the supplement.

2.5 Centering test statistics

Our refined statistic $\eta_t^s(\hat{g}_n)$ needs to be further centered to ensure asymptotic normality. To achieve this, ideally we adjust \hat{g}_n in the direction of the score function $\alpha_{t,n}^*$ to attain a smaller loss than in (7). With the estimate $\hat{\alpha}_n$ of $\alpha_{t,n}^*$ given by (13), one naturally uses

$$\check{g}_n = \hat{g}_n + \hat{\delta}_{\mathrm{t}} \hat{\alpha}_n \tag{14}$$

to arrive at a debiased statistic $\eta_t^{s}(\check{g}_n)$, where

$$\widehat{\delta}_{t} = \frac{1}{\sum_{i=1}^{n} \widehat{\alpha}_{n}^{2}(\widetilde{F}_{i}, X_{i,j})} \sum_{i=1}^{n} \left\{ Y_{i} - \widehat{g}_{n}(\widetilde{F}_{i}, X_{i,j}) \right\} \widehat{\alpha}_{n}(\widetilde{F}_{i}, X_{i,j})$$
(15)

is chosen to minimize $\mathbb{P}_n \ell(\cdot; \hat{g}_n + \delta_t \hat{\alpha}_n)$, as derived in Section C.3. We can interpret this update step as a form of the targeted machine learning (van der Laan and Rose, 2018)

specialized to our deep neural network context. In Section 3.3 we will show that \check{g}_n indeed satisfies an approximate minimization condition: for an infinitesimal stepsize $r_{\inf,n}$,

$$\mathbb{P}_n\ell(\cdot;\check{g}_n) - \mathbb{P}_n\ell(\cdot;\check{g}_n \pm r_{\inf,n}\alpha^*_{t,n}) \leqslant r_{\inf,n}b_n \tag{16}$$

for a small tolerance b_n . This will in turn center our test statistics.

2.6 Final test statistic and its uniform extensions

To alleviate the technical possibility of the unboundnessed of our derivative estimator, we let $\Psi : \mathbb{R} \to \mathbb{R}$ be a bounded truncation function that also satisfies Assumption B.5; we provide an example of such a truncation function below the assumption. We then define the final (fixed-t) test statistic to be $\check{\eta}_t^{s}(\check{g}_n)$ where the operator $\check{\eta}_t^{s}(\cdot)$ acts on functions $g: \mathbb{R}^{\bar{r}+1} \to \mathbb{R}$ as

$$\check{\eta}_t^{\mathrm{s}}(g) = \frac{1}{n} \sum_{i \in \mathcal{I}_h} \left[\exp\{t \, \Psi(g_j^{\mathrm{s}}(\widetilde{F}_i, X_{i,j}))\} - 1 \right].$$
(17)

Its standard error is estimated as $t \| \hat{\epsilon} \hat{\alpha}_n \|_{L_2(\mathbb{P}_n)} = t [\sum_{i=1}^n \{ \hat{\epsilon}_i \hat{\alpha}_n (\tilde{F}_i, X_{i,j}) \}^2 / n]^{1/2}$, where $\hat{\epsilon}_i = Y_i - \check{g}_n (\tilde{F}_i, X_i)$ is the *i*-th residual. This leads to a level- α test of the null hypothesis (3) with the critical region

$$\left|\frac{\sqrt{n}}{t}\frac{\check{\eta}_t^{\mathrm{s}}(\check{g}_n)}{\|\epsilon\alpha_n^{\mathrm{null}}\|_{L_2}}\right| > z_{1-\alpha/2} \tag{18}$$

where $z_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of a standard normal distribution; the choice of this critical value is supported by the first part of (33) in Theorem 3.5.

To increase the power of our test, we aggregate the ensemble of test statistics $\check{\eta}_t^s(\check{g}_n)$ over a range of t. We consider two types of aggregation: a sup-statistic and an integrated weighted square-statistic (or simply a square-statistic). In both cases, let $\delta > 0$ be a small but fixed constant, let $T > \delta$ be a potentially large but also fixed constant, and define $\mathcal{T}_{\delta} = [-T, -\delta] \cup [\delta, T]$. The two aforementioned aggregations of test statistics and their associated critical regions are as follows. 1) Sup-statistic:

$$\widehat{Z} = \sqrt{n} \sup_{t \in \mathcal{T}_{\delta}} \left| \frac{\widecheck{\eta}_t^{\mathbf{s}}(\widecheck{g}_n)}{t} \right|, \quad \text{with critical region: } \frac{\widehat{Z}}{\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}} > z_{1-\alpha/2}.$$
(19)

2) Square statistic:

$$\widehat{\chi}^2 = \frac{n}{\int_{\mathcal{T}_{\delta}} w(t) \mathrm{d}t} \int_{\mathcal{T}_{\delta}} \frac{1}{t^2} \left\{ \widecheck{\eta}_t^{\mathrm{s}}(\widecheck{g}_n) \right\}^2 w(t) \mathrm{d}t, \quad \text{with critical region } \frac{\widehat{\chi}^2}{\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}^2} > \chi_{1,1-\alpha}^2$$
(20)

where w is an integrable weight function such as $w(t) = \exp(-\beta t^2)$ for a given $\beta > 0$, and $\chi^2_{1,1-\alpha}$ is the $1-\alpha$ quantile of a chi-square distribution with one degree of freedom.

The critical values are supported by the last two parts of (33) in Theorem 3.5.

2.7 Summary of proposed methods

The major implementation steps of our partial derivative estimation and conditional screening test procedures in Section 2 are summarized in Algorithm 2.1.

3 Theoretical properties of estimators and tests

3.1 Initial regression and smoothed derivative estimators

We first collect the assumptions most relevant to the construction and rates of our estimators \hat{g}_n and $\hat{g}_{n,j}^s$; we defer the more technical assumptions to Section B.1. In particular, Assumption 3.1 imposes mild conditions on the data generating processes, and Assumption 3.2 specifies the function class for m_0 and its associated neural networks.

Assumption 3.1 (Data generating processes). (i) (\mathbf{F}, X_j) take value on the bounded support $[-b, b]^r \times [-b, b]$; (ii) The dimension \overline{r} of the diversified factor $\widetilde{\mathbf{F}}$ is a fixed constant and satisfies $\overline{r} \ge r$; (iii) Regression function $m_0^*(\mathbf{f}, \mathbf{x})$, defined below (4), is uniformly bounded in magnitude by M_∞ that satisfies $M_\infty \le M$ for M the truncation level in (6).

- **Input:** Observed samples $Y_i, \mathbf{X}_i \in \mathbb{R}^d, i \in \{1, \dots, n\}$, and an exogenous diversified projection matrix $\mathbf{W} \in \mathbb{R}^{d \times \overline{r}}$; specification of the coordinate j for screening.
- **Result:** Preliminary and refined regression estimators $\hat{g}_n(\cdot)$ and $\check{g}_n(\cdot)$; their
 - smoothed derivative estimators $\hat{g}_{n,j}^{s}(\cdot)$ and $\check{g}_{n,j}^{s}(\cdot)$; $\hat{\alpha}_{n}(\cdot)$, and test results.
 - 1. Compute $\widetilde{F}_i = d^{-1} \mathbf{W}^\top \mathbf{X}_i$, for all $i \in \{1, \ldots, n\}$.
 - Estimate the regression function m₀ in (2) by the neural network estimator ĝ_n from (7). Conduct cross validation (see Remark A.4), or otherwise, to find a smoothing bandwidth h.
 - 3. Define the loss function $\widehat{R}_n^{\text{null}}(\cdot)$ as

$$\widehat{R}_{n}^{\text{null}}(\alpha) = \frac{1}{n} \sum_{i=1}^{n} \alpha^{2}(\widetilde{F}_{i}, X_{i,j}) - 2\frac{1}{n} \sum_{i \in \mathcal{I}_{h}} \alpha_{j}^{\text{s}}(\widetilde{F}_{i}, X_{i,j})$$

which can be approximated via Lemma A.2, and let the estimator $\hat{\alpha}_n$ of $\alpha_{t,n}^*$ be

$$\widehat{\alpha}_n = \operatorname{argmin}_{\alpha \in \mathcal{F}_n(\overline{r}+1)} \widehat{R}_n^{\operatorname{null}}(\alpha).$$

4. Return $\check{g}_n = \hat{g}_n + \hat{\delta}_t \hat{\alpha}_n$, where

$$\widehat{\delta}_{t} = \frac{1}{\frac{1}{n}\sum_{i=1}^{n}\widehat{\alpha}_{n}^{2}(\widetilde{F}_{i}, X_{i,j})} \frac{1}{n} \sum_{i=1}^{n} \left\{ Y_{i} - \widehat{g}_{n}(\widetilde{F}_{i}, X_{i,j}) \right\} \widehat{\alpha}_{n}(\widetilde{F}_{i}, X_{i,j})$$

- 5. Let $\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}$ be the standard deviation estimator of our test statistic $\check{\eta}_t^{\mathrm{s}}(\check{g}_n)$ (with $\check{\eta}_t^{\mathrm{s}}$ from (17)); specifically, $\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)} = [\frac{1}{n}\sum_{i=1}^n \{\widehat{\epsilon}_i\widehat{\alpha}_n(\widetilde{F}_i, X_{i,j})\}^2]^{1/2}$ where $\widehat{\epsilon}_i = Y_i \check{g}_n(\widetilde{F}_i, X_i)$ is the *i*-th residual. Then, for a prescribed significance level $0 < \alpha < 1$, reject the null hypothesis (3) in our conditional screening hypotheses:
 - a) For the fixed-t test, if (18) holds;
 - b) For the sup test, if the decision rule in (19) holds;
 - c) For the square test, if the decision rule in (20) hold.

Algorithm 2.1: Partial derivative estimation & conditional screening test by deep neural network

Assumption 3.2 (Function class for m_0 and neural network scaling). Regression function m_0 belongs to the hierarchical composition model $\mathcal{H}(r + 1, l, \mathcal{P}, C_{\mathcal{H}})$ in Definition A.5 (Schmidt-Hieber, 2020; Kohler and Langer, 2021). The tuple $(L, \mathbf{k} = (\bar{r} + 1, k_1, \dots, k_L, 1))$ for the structure of $\mathcal{F}_n(\bar{r} + 1)$ in (6) satisfies $k_1 = \cdots = k_L \equiv k_0$ for some $k_0 \in \mathbb{N}$ and $L \cdot k_0 \sim n^{\frac{1}{4\kappa+2}} \log^{\frac{4\kappa-1}{2\kappa+1}}(n)$ where κ is the dimension-adjusted degree of smoothness defined above Definition A.5.

Assumption 3.3 (Idiosyncratic terms). The idiosyncratic terms $\boldsymbol{u} \in \mathbb{R}^d$ satisfy $\sum_{j=1}^d \mathbb{E}[u_j^2] \lesssim d_u$ for $d_u \leq d$ and a weak dependence condition $\sum_{j \neq j'} |\mathbb{E}[u_j u_{j'}]| \lesssim d_u$.

In Theorem 3.1 we establish the convergence of $\hat{g}_n - m_0$ for regression function estimation and of $\hat{g}_{n,j}^s - m_{0,j}$ for derivative estimation. The following rates will appear:

$$p_n = c_p n^{\frac{1}{2\kappa+1}} \log^{2\frac{4\kappa-1}{2\kappa+1}+2}(n), \quad \nu_n = n^{-\frac{\kappa}{2\kappa+1}} \log^{\frac{6\kappa}{2\kappa+1}}(n), \quad \delta_f = \{\overline{r}d_u/d^2\}^{1/2}$$
(21)

where $c_{\rm p}$ is the constant appearing in (C.9), d_u appears in Assumption 3.3, and $\delta_{\rm f} \to 0$ as $d \to \infty$. Moreover, because the target of the *smoothed* derivative estimator $\hat{g}_{n,j}^{\rm s}$ is the *smoothed* derivative $m_{0,j}^{\rm s}$, not $m_{0,j}$, naturally a bias

$$r_{\mathbf{b},m,j} \equiv \{\int_{\Omega_h} (m_{0,j}^{\mathbf{s}} - m_{0,j})^2 \,\mathrm{d}\mathbb{P}\}^{1/2} = \left\{\int_{\Omega_h} [\int_{-h}^{h} \{m_{0,j}(\boldsymbol{f}, x_j - z) - m_{0,j}(\boldsymbol{f}, x_j)\} K_h(z) \mathrm{d}z]^2 \mathrm{d}\mathbb{P}\right\}^{1/2}$$

occurs. However, this bias vanishes under the null in (3) because there $m_{0,j} = 0$ everywhere.

Theorem 3.1. Under Assumptions 3.1 to 3.3 and B.1 to B.3, for p_n , ν_n and δ_f in (21), on an event \mathcal{A}'_m with $\mathbb{P}(\mathcal{A}'_m) \ge 1 - C\exp(-p_n)$, and for a constant $c'_{m,1}$, the initial regression function estimator \hat{g}_n from (7) satisfies

$$\left[\int \{ \widehat{g}_n(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, x_j) \}^2 \mathrm{d}\mathbb{P} \right]^{1/2} \leqslant c'_{m,1}(\nu_n + \delta_\mathrm{f}) \equiv r_{m,n}.$$
(22)

In addition, its associated smoothed derivative estimator $\hat{g}_{n,j}^{s}$ satisfies, on the same event \mathcal{A}'_{m} and with a possibly different constant $c'_{m,2}$,

$$\left[\int_{\Omega_h} \{\widehat{g}_{n,j}^{s}(\widetilde{f}, x_j) - m_{0,j}(f, x_j)\}^2 d\mathbb{P}\right]^{1/2} \leqslant c'_{m,2} h^{-1}(\nu_n + \delta_f) + r_{b,m,j}.$$
(23)

The proof of Theorem 3.1 is deferred to Section C.1. With the rate $r_{m,n}$ in (22) for the regression estimator, the rate (23) for the smoothed *derivative* estimator follows from Lemma A.4 that applies to the smoothed derivatives of all estimators. Next, the rate $r_{m,n}$ consists of two parts: ν_n represents the combined effect of stochastic error and neural network approximation bias, while δ_f represents the error induced by relying on the diversified factors in place of the latent factors. Theorem 3.1 also covers as a special example the low-dimensional \boldsymbol{X} case in Section D.1 by simply setting $\delta_f = 0$.

3.2 Estimating the asymptotic variance

3.2.1 Definition of the score function In this section, to complete our discussion in Section 2.4, we formally introduce the score function $\alpha_{t,n}^*$ of our test statistics and the bias of $\hat{\alpha}_n$ relative to $\alpha_{t,n}^*$. Define the population version of the operator $\check{\eta}_t^{s}(g)$ in (17) as

$$\widetilde{\eta}_t^{s}(g) = \int_{\Omega_h} \left[\exp\{t \, \Psi(g_j^{s}(\widetilde{f}, x_j))\} - 1 \right] d\mathbb{P},$$
(24)

and its associated directional derivative (in the direction $v = v(\tilde{f}, x_j)$)

$$\frac{\partial \tilde{\eta}_t^{s}(g_0)}{\partial g} [v] = \left. \frac{\partial \tilde{\eta}_t^{s}(g_0 + \tau v)}{\partial \tau} \right|_{\tau=0}$$

By (A.3) in the proof of Lemma C.4, $\frac{\partial \tilde{\eta}_t^s(g_0)}{\partial g}[v]/t$ is a bounded linear functional of v. Hence, there exists a *Riesz representer* that becomes our $\alpha_{t,n}^*$ and that satisfies $\frac{\partial \tilde{\eta}_t(g_0)}{\partial g}[v]/t = \int v \alpha_{t,n}^* d\mathbb{P}$ for all v (Conway, 1990, Theorem 3.4, Chapter 1). In particular, by (A.3), the effect of $\alpha_{t,n}^*$ is captured explicitly as

$$\int v \alpha_{t,n}^* d\mathbb{P} = \frac{1}{t} \frac{\partial \widetilde{\eta}_t^{\mathrm{s}}(g_0)}{\partial g} [v] = \int_{\Omega_h} \exp\{t \,\Psi(g_{0,j}^{\mathrm{s}}(\widetilde{f}, x_j))\} \dot{\Psi}\{g_{0,j}^{\mathrm{s}}(\widetilde{f}, x_j)\} v_j^{\mathrm{s}}(\widetilde{f}, x_j) d\mathbb{P}$$

$$\overset{\text{under } H_0}{=} \int_{\Omega_h} v_j^{\mathrm{s}}(\widetilde{f}, x_j) d\mathbb{P}.$$
(25)

Because the loss \hat{R}_n^{null} relies on the last step of (25), $\hat{\alpha}_n$ consistently estimates $\alpha_{t,n}^*$ under the null, as confirmed by Theorem 3.2. However $\hat{\alpha}_n$ may not be consistent for $\alpha_{t,n}^*$ under the alternative. Instead, it is not hard to show that it is consistent for a population limit α_n^{null} which exists and satisfies $\int v \alpha_n^{\text{null}} d\mathbb{P} = \int_{\Omega_h} v_j^{\text{s}}(\tilde{f}, x_j) d\mathbb{P}$ for all v. See the remark below Eq. (A.3) for details.

To obtain an intuitive idea of what the Riesz representer $\alpha_{t,n}^*$ could look like, assume that \tilde{F}, X_j admit a joint density $p(\tilde{f}, x_j)$ that is differentiable in x_j with the derivative being $p_j(\tilde{f}, x_j)$. Then, by Lemma A.3, under the null hypothesis and in the limit $h \to 0$, $\alpha_{t,n}^*(\tilde{f}, x_j) = -p_j(\tilde{f}, x_j)/p(\tilde{f}, x_j)$; in this case, if further \tilde{F}, X_j are jointly Gaussian on their support, then $\alpha_{t,n}^*(\tilde{f}, x_j)$ is simply a linear function in \tilde{f} and x_j .

3.2.2 Convergence rate for the score function estimator To ensure the rate of $\hat{\alpha}_n$, our Assumption 3.4 mirrors our earlier Assumptions 3.1 and 3.2 for estimating the regression function m_0 with deep neural networks and in particular imposes a hierarchical composition model on α_n^{null} that takes the diversified predictors (\tilde{F}, X_j) as arguments. Then, Assumption 3.5 places a mild condition on the smoothing bandwidth. As in Section 3.1, we defer the more technical assumptions to Section B.1.

Assumption 3.4 (Function class and neural network scaling for estimating $\alpha_{t,n}^*$). The function α_n^{null} , when we restrict the support to $[-c_b b, c_b b]^{\overline{r}} \times [-b, b]$ for a constant $c_b > 0$, belongs to the hierarchical composition model $\mathcal{H}(\overline{r}+1, l, \mathcal{P}, C_{\mathcal{H}})$ on the same support and (without loss of generality) with the same $l, \mathcal{P}, C_{\mathcal{H}}$ as in Assumption 3.2. The class $\mathcal{F}_n(\overline{r}+1)$ in (13) satisfies the same structural scaling as in Assumption 3.2.

Assumption 3.5 (Rate of bandwidth). The bandwidth h satisfies $h \ge 1/(\sqrt{n\nu_n})$.

To characterize the bias of α_n^{null} with respect to $\alpha_{t,n}^*$ under the alternative, define the random variable that represents the signal of the alternative as

$$Z_{t,j,h} = \left[\exp(t\,\Psi\{g_{0,j}^{\mathbf{s}}(\widetilde{\boldsymbol{F}},X_j)\}) - 1\right]\mathbb{1}\{X_j \in \mathcal{B}_h\}.$$
(26)

Note that $Z_{t,j,h}$ always equals zero under the null hypothesis.

Theorem 3.2. Under Assumptions 3.1 to 3.5 and B.1 to B.4, for p_n , ν_n and δ_f in (21), on an event \mathcal{A}_{α} with $\mathbb{P}(\mathcal{A}_{\alpha}) \ge 1 - C \exp(-p_n)$, and for a constant $c_{\alpha,1}$,

$$\left[\int (\widehat{\alpha}_n - \alpha_n^{\text{null}})^2 (\widetilde{\boldsymbol{f}}, x_j) d\mathbb{P}\right]^{1/2} \leqslant c_{\alpha, 1} \{ (h^{-1}\nu_n + \delta_f) \wedge M \} \equiv r_{\alpha, n}^{\text{null}}.$$
 (27)

Moreover, for a constant $c_{\alpha,2}$ not dependent on t, the bias of α_n^{null} relative to $\alpha_{t,n}^*$ is

$$\forall t \in \mathbb{R}, \quad \left[\int (\alpha_n^{\text{null}} - \alpha_{t,n}^*)^2 (\widetilde{\boldsymbol{f}}, x_j) d\mathbb{P} \right]^{1/2} \leqslant c_{\alpha,2} (h^{-1} \| Z_{t,j,h} \|_{L_2} \wedge M) \equiv r_{\alpha,t,b}. \tag{28}$$

Consequently, on the event \mathcal{A}_{α} ,

$$\forall t \in \mathbb{R}, \quad \left[\int (\hat{\alpha}_n - \alpha_{t,n}^*)^2 (\widetilde{f}, x_j) \mathrm{d}\mathbb{P} \right]^{1/2} \leqslant r_{\alpha,n}^{\mathrm{null}} + r_{\alpha,t,\mathrm{b}} \equiv r_{\alpha,t,n}. \tag{29}$$

The proof of Theorem 3.2 is deferred to Section C.2. Compared with the rate (22) in Theorem 3.1 for the regression estimator \hat{g}_n , the rates for the score function estimator $\hat{\alpha}_n$ mainly differ in two aspects: first, a factor h^{-1} precedes ν_n which is the consequence of the smoothing operation in the loss function \hat{R}_n^{null} in (12); second, under the alternative hypothesis a bias $h^{-1} ||Z_{t,j,h}||_{L_2}$ in (28) is induced relative to $\alpha_{t,n}^*$.

3.3 Centering and adjusted estimator

Recall from Section 2.5 that the adjusted estimator \check{g}_n in (14) was introduced to center our test statistics through a suitable minimization condition (16). In this section Proposition 3.3 first shows that a small $\hat{\delta}_t$ is sufficient to arrive at \check{g}_n that satisfies (16) with a small b_n .

Assumption 3.6 (Centering test statistics). *n* is large enough such that (i) $c_{t,1}(\delta_f + \nu_n/h) \leq 1$ for a large enough constant $c_{t,1}$; (ii) $c_{t,2}r_{m,n} \leq 1$ for the constant $c_{t,2}$ in Proposition 3.3; also, let the infinitesimal positive sequence $r_{inf,n}$ satisfy $r_{inf,n} = \mathcal{O}(r_{m,n}\nu_n)$.

Proposition 3.3. Under Assumptions 3.1 to 3.6(i) and B.1 to B.4, on an event $\mathcal{A}_{t,1}$ with $\mathbb{P}(\mathcal{A}_{t,1}) \ge 1 - C \exp(-p_n)$, for a constant $c_{t,2}$, $\hat{\delta}_t$ given by (15) satisfies $|\hat{\delta}_t| \le c_{t,2}r_{m,n}$. If

furthermore Assumption 3.6(ii) holds, then on an event $\mathcal{A}_{t,2}$ with $\mathbb{P}(\mathcal{A}_{t,2}) \ge 1 - C \exp(-p_n)$, for a constant $c_{t,3}$, condition (16) holds uniformly at all $t \in \mathcal{T}_{\delta}$ and $b_n = c_{t,3}r_{m,n}(r_{\alpha,n}^{null} + r_{\alpha,t,b})$.

The proof of Proposition 3.3 is deferred to Section C.3. In the proposition, under the null the tolerance $b_n = r_{m,n}r_{\alpha,n}^{\text{null}}$, which is faster than $n^{-1/2}$ under appropriate conditions and will ensure the asymptotic normality of our test statistics. Under the alternative, b_n is not necessarily faster than $n^{-1/2}$, but is still fast enough to ensure consistency under the local alternatives.

We will start working mostly with the adjusted estimator \check{g}_n from now on. The next theorem is the counterpart of Theorem 3.1 and shows that \check{g}_n and its smoothed derivative estimator $\check{g}_{n,j}^s$ maintain convergence rates similar to the unadjusted ones.

Theorem 3.4. Under Assumptions 3.1 to 3.6 and B.1 to B.4, for p_n , ν_n and δ_f in (21), on an event \mathcal{A}_m with $\mathbb{P}(\mathcal{A}_m) \ge 1 - C \exp(-p_n)$, and for a constant $c_{m,1}$,

$$\left[\int \{ \check{g}_n(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, \boldsymbol{x}) \}^2 \mathrm{d}\mathbb{P} \right]^{1/2} \leqslant c_{m,1}(\nu_n + \delta_\mathrm{f}).$$
(30)

In addition, its smoothed derivative estimator $\check{g}_{n,j}^{s}$ satisfies, on the same event \mathcal{A}_{m} and with a possibly different constant $c_{m,2}$,

$$\left[\int_{\Omega_h} \{\widetilde{g}_{n,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j) - m_{0,j}(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P}\right]^{1/2} \leqslant c_{m,2}h^{-1}(\nu_n + \delta_{\mathrm{f}}) + r_{\mathrm{b},m,j}.$$
(31)

The proof of Theorem 3.4 is deferred to Section C.4.

3.4 Properties of the conditional screening tests

This section gives the result on the asymptotic null distributions of the test statistics. The results for consistency against the local alternatives are given in Section C.6.

Theorem 3.5. Suppose that Assumptions 3.1 to 3.6 and B.1 to B.4 hold, and in addition condition (*) holds: $\delta_{\rm f} + h^{-2}(\nu_n^2 + \delta_{\rm f}^2) = \mathcal{O}(n^{-1/2})$. Then, with \mathcal{T}_{δ} and tests given in Section 2.6, we have

for all fixed
$$t \in \mathcal{T}_{\delta}, \frac{\sqrt{n}}{t} \frac{\check{\eta}_{t}^{s}(\check{g}_{n})}{\|\epsilon\alpha_{n}^{null}\|_{L_{2}}} \rightarrow_{d} \mathcal{Z}, \quad \frac{\widehat{Z}}{\|\epsilon\alpha_{n}^{null}\|_{L_{2}}} \rightarrow_{d} |\mathcal{Z}|, \quad \frac{\widehat{\chi}^{2}}{\|\epsilon\alpha_{n}^{null}\|_{L_{2}}^{2}} \rightarrow_{d} \chi_{1}^{2}, \quad (32)$$

where \mathcal{Z} stands for a standard normal random variable and χ_1^2 is a chi-square random variable with one degree of freedom. Moreover, we are free to replace $\|\epsilon \alpha_n^{\text{null}}\|_{L_2}$ in (32) above by its empirical counterpart $\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}$ from Section 2.6, to conclude that

for all fixed
$$t \in \mathcal{T}_{\delta}, \frac{\sqrt{n}}{t} \frac{\check{\eta}_{t}^{s}(\check{g}_{n})}{\|\widehat{\epsilon}\widehat{\alpha}_{n}\|_{L_{2}(\mathbb{P}_{n})}} \to_{d} \mathcal{Z}, \quad \frac{\widehat{Z}}{\|\widehat{\epsilon}\widehat{\alpha}_{n}\|_{L_{2}(\mathbb{P}_{n})}} \to_{d} |\mathcal{Z}|, \quad \frac{\widehat{\chi}^{2}}{\|\widehat{\epsilon}\widehat{\alpha}_{n}\|_{L_{2}(\mathbb{P}_{n})}^{2}} \to_{d} \chi_{1}^{2}.$$
 (33)

The proof of Theorem 3.5 is deferred to Section C.5. The extra condition (*) in Theorem 3.5 is natural due to the presence of the \sqrt{n} scaling factor in the our tests and is mild. Note that if the bandwidth h is held as a constant, then condition (*) simply reduces to $\delta_{\rm f} = \mathcal{O}(n^{-1/2})$, and $\nu_n = \mathcal{O}(n^{-1/4})$ which is in turn implied by $\kappa > 1/2$ for κ in Assumption 3.2 on the hierarchical composition model.

4 Simulation studies

In this section, we conduct simulations to assess the size and power of our conditional screening tests across different scenarios. Recall that we summarize the implementations of these tests in Algorithm 2.1. We consistently employ the Quartic/biweight kernel.

We choose t = 1 for the fixed-t test statistic, and simply set $\mathcal{T}_{\delta} = \{-1.25, -0.5, 0.5, 1.25\}$ in the sup statistic and the square statistic in (19) and (20) respectively. Then, we implement the sup statistic in (19) as $\sup_{t \in T_{\delta}} |\hat{\eta}_t^s(\check{g}_n)|/t$, and the square statistic as $\hat{\chi}_{\delta}^2 = n \sum_{t \in \mathcal{T}_{\delta}} {\{\hat{\eta}_t^s(\check{g}_n)\}^2/t^2}$ (we take w(t) = 1 in (20)).

We set d = 200 and d = 400 when n = 256 and n = 512 respectively, and the full regression model as $Y = m_0^*(\mathbf{F}, X_1, X_3) + \epsilon^*$ (see m_0^* defined below (4)). To illustrate the performance of our conditional screening test, we have designed m_0^* so that a significant portion of its variation is accounted for by the factors. We further consider a nonlinear and a linear model of m_0^* . Specifically, under the null hypothesis, we set:

nonlinear :
$$m_0^*(\cdot) = m_0^{*(null)}(\cdot) = \sin(f_1 + u_1) + \log(8 + f_2) \times \log(8 + f_3) + \exp(-f_4^2/2),$$

linear : $m_0^*(\cdot) = m_0^{*(null)}(\cdot) = f_1 - f_2 + f_3 + f_4 - f_5.$
(34)

For our screening test, we always select X_3 as the variable of interest. Accordingly, under the alternative hypothesis we add signals in X_3 to $m_0^{*(null)}$ above:

nonlinear :
$$m_0^*(\cdot) = m_0^{*(null)}(\cdot) + X_3^2/4$$
, linear : $m_0^*(\cdot) = m_0^{*(null)}(\cdot) + X_3/16$. (35)

The signal under the linear model is very weak and allows us to discern different test settings. While the signal under the nonlinear model may not appear weak at first, here the *average* derivative $\mathbb{E}m_{0,j} = 0$, so detecting departure from the null critically depends on the higher-order, nonlinear effect of our MGF test statistics. Under the nonlinear model r = 4 while under the linear model r = 5. In both models we set $\bar{r} = r$.

Under the nonlinear model, u_1 is incorporated for a richer structure that specifically leads to the working regression function $m_0(\mathbf{F}, X_3)$ being different from the full regression function $m_0^*(\mathbf{F}, X_1, X_3)$ by design; see Remark A.2. We incorporate $u_1 = X_1 - \mathbf{B}_1 \cdot \mathbf{F}$ instead of X_1 directly because u_1 will be enforced to be independent of X_3 under screening. Moreover, under the null hypothesis, the selection of X_3 is purely for clarity: under the nonlinear model, we could conduct our test on any X_j for any $j \in \{2, \ldots, d\}$, and under the linear model, for any $j \in \{1, \ldots, d\}$.

For the high-dimensional predictor X in the factor model (1), the factor loading matrix **B** is generated with i.i.d. Unif $[-\sqrt{3}, \sqrt{3}]$ entries; the factor F and the idiosyncratic terms u both have i.i.d. $\mathcal{N}(0, 0.6)$ entries. We further draw the noise as $\epsilon^* \sim \mathcal{N}(0, 0.3)$. The quantities **B**, F, u and ϵ^* are all drawn independently. We pre-train **W** with samples of size 100.

We set the hyper-parameters for the neural network fitting as follows: for the regression estimator \hat{g}_n in (7), we employ neural networks with L = 5 hidden layers, a common width of $k_0 = 16$ per layer, and the ReLU activation function. Training proceeds over 800 epochs, utilizing a batch size of 256 and a constant learning rate of 0.005. For the score function estimator $\hat{\alpha}_n$ in (13), we maintain the same neural network fitting parameters, except that we set L = 2, batch size to 64 and the number of epochs to 400. We employ early stopping as the only regularization technique, and terminate training if there's no improvement after 20 epochs (patience = 20) on a validation set.

We conduct our tests at significance levels of either 5% or 10% under both the nonlinear and the linear models, resulting in a total of four combinations summarized in Tables 1 to 4. In each table we present in alternating rows the sizes and powers (the latters in parentheses) of the tests in (18), (19) and (20), and of the same tests but without centering (that is, tests employing the non-adjusted estimator \hat{g}_n), under different sample sizes. The columns, arranged from left to right, correspond to the four different bandwidths. Each entry is calculated based on 500 Monte Carlo repetitions.

We make a few observations from the results in the tables. First, centering generally improves the size of our tests without affecting their power. Moreover, under the null, sizes arrive at the nominal level if we further increase the bandwidth; this is especially evident for the sup test, which naturally tends to reject more often (under both the null and alternative). This shows the improvement of the power of the centered, smoothed test compared to the initial test (9) if both tests are calibrated to have the same size. Of course, if we increase h more than 1.5, we expect the biases will arrive under the alternative. In the nonlinear case (Tables 1 and 3), the ensemble tests by aggregating various t also outperform the fixed-t test in power.

		h	0.1	1.0	1.5	2.0
	fixed- t test	Non-centered	0.16 (0.67)	0.14 (0.57)	$0.13 \ (0.56)$	0.12(0.51)
n = 256		Centered	$0.11 \ (0.69)$	$0.10 \ (0.60)$	0.08~(0.55)	0.07~(0.50)
	sup test	Non-centered	$0.26\ (0.97)$	0.21 (0.93)	0.19(0.92)	0.18(0.87)
		Centered	$0.19\ (0.97)$	0.16(0.93)	$0.13\ (0.91)$	$0.13 \ (0.85)$
	square test	Non-centered	$0.16 \ (0.86)$	0.14 (0.79)	0.14(0.74)	$0.12 \ (0.65)$
		Centered	$0.11 \ (0.82)$	$0.09 \ (0.78)$	$0.09 \ (0.70)$	$0.09 \ (0.63)$
	fixed- t test	Non-centered	$0.11 \ (0.88)$	0.10 (0.80)	$0.10 \ (0.77)$	0.08~(0.70)
		Centered	$0.09 \ (0.88)$	$0.07 \ (0.83)$	$0.07 \ (0.76)$	0.05~(0.70)
n = 512	sup test	Non-centered	0.19 (1.00)	0.16 (0.98)	$0.15 \ (0.97)$	$0.11 \ (0.96)$
		Centered	0.15(1.00)	$0.12 \ (0.97)$	$0.11 \ (0.97)$	0.08~(0.95)
	square test	Non-centered	$0.12 \ (0.96)$	0.10 (0.91)	0.09(0.87)	0.08 (0.85)
		Centered	0.08~(0.96)	$0.07 \ (0.91)$	0.05~(0.87)	0.05~(0.86)

Table 1: Performance summary of our test statistics under the nonlinear model in (34) (for size under the null) and (35) (for power under the alternative) at the significance level 5%. Specifically, we provide the size and power (the latter displayed in parentheses) of our tests under various combinations of the test statistic (fixed-t, sup, or squared statistic), sample size, bandwidth, and the use of either the non-centered (employing the non-adjusted estimator \hat{g}_n) or the centered (employing the adjusted \check{g}_n) test statistics. Each value in the table represents the average over 500 Monte Carlo repetitions.

		h	0.1	1.0	1.5	2.0
		Non-centered	0.16 (0.84)	0.14 (0.78)	0.14 (0.76)	0.14(0.76)
	inxed - ι test	Centered	0.08~(0.87)	0.09(0.82)	0.08~(0.80)	$0.07 \ (0.79)$
n = 512	sup test	Non-centered	$0.25 \ (0.85)$	$0.21 \ (0.79)$	0.18(0.77)	$0.18 \ (0.76)$
		Centered	0.18(0.88)	$0.13 \ (0.83)$	$0.11 \ (0.80)$	$0.11 \ (0.80)$
	square test	Non-centered	0.14(0.74)	0.14(0.72)	$0.15 \ (0.72)$	0.14(0.72)
		Centered	$0.09 \ (0.77)$	0.09~(0.75)	0.08~(0.73)	$0.07 \ (0.75)$

Table 2: Performance summary of our test statistics under the linear model in (34) (for size under the null) and (35) (for power under the alternative) at the significance level 5%.

		h	0.1	1.0	1.5	2.0
	fixed- t test	Non-centered	$0.17 \ (0.90)$	0.18 (0.84)	0.18(0.81)	0.14(0.76)
n = 512		Centered	$0.13\ (0.91)$	$0.13 \ (0.86)$	0.13(0.81)	$0.13\ (0.77)$
	sup test	Non-centered	0.27(1.00)	0.24 (0.99)	0.24(0.98)	$0.20 \ (0.97)$
		Centered	0.24(1.00)	$0.19 \ (0.99)$	$0.19 \ (0.98)$	$0.17 \ (0.98)$
	square test	Non-centered	0.19(0.99)	0.18 (0.96)	0.16(0.93)	$0.13\ (0.90)$
		Centered	$0.13 \ (0.98)$	$0.12 \ (0.95)$	$0.12 \ (0.92)$	$0.11 \ (0.90)$

Table 3: The same caption as Table 1 except for significant level 10%.

		h	0.1	1.0	1.5	2.0
		Non-centered	$0.23 \ (0.89)$	$0.21 \ (0.85)$	0.20(0.82)	$0.21 \ (0.82)$
	$\pi xed-t$ test	Centered	0.14(0.93)	$0.16 \ (0.88)$	$0.13 \ (0.87)$	$0.12 \ (0.85)$
n = 512	sup test	Non-centered	0.34(0.90)	$0.26 \ (0.86)$	0.26(0.82)	0.26(0.83)
		Centered	$0.26\ (0.93)$	$0.21 \ (0.89)$	$0.17 \ (0.88)$	$0.16\ (0.86)$
	square test	Non-centered	0.21 (0.83)	0.21 (0.80)	$0.20 \ (0.78)$	$0.21 \ (0.79)$
		Centered	$0.15 \ (0.86)$	0.15(0.84)	0.12(0.83)	0.13(0.82)

Table 4: The same caption as Table 2 except for significant level 10%.

5 Empirical Applications

5.1 Asset Pricing

In our first empirical analysis, we examine a comprehensive dataset containing both returns and specific characteristics of firms. Twelve monthly returns in 2021 for the 100 largest financial institutions are drawn from the CRSP database as our response variable, along with a set of d = 48 firm-specific characteristics as predictors; the resulting sample size is n = 1200. The dataset's foundation is credited to Chen and Zimmermann (2022).

We select the smoothing bandwidth from the sequence $\{0.4, 0.6, \dots, 1.8, 2.0\}$ using the bandwidth selection algorithm suggested in Remark A.4. We set $\overline{r} = 5$ the conventional practice in finance, the hidden size $k_0 = 24$, and the learning rate $\tau = 10^{-3}$; we keep the other neural network fitting parameters identical to Section 4.

In Figure 1 we plot for each of the 44 continuous predictors the intervals whose right ends are the test statistics normalized by their corresponding critical values in (18), (19) and (20). Specifically, we plot the intervals

$$[0, n^{1/2} | \tilde{\eta}_t^{\mathsf{s}}(\tilde{g}_n) | / (tz_{1-\alpha/2} \| \hat{\epsilon} \hat{\alpha}_n \|_{L_2(\mathbb{P}_n)})], \quad [0, \hat{Z} / (z_{1-\alpha/2} \| \hat{\epsilon} \hat{\alpha}_n \|_{L_2(\mathbb{P}_n)})], \quad [0, \sqrt{\hat{\chi}^2} / (\sqrt{\chi_{1,1-\alpha}^2} \| \hat{\epsilon} \hat{\alpha}_n \|_{L_2(\mathbb{P}_n)})]$$

for the fixed-*t*, sup, and square tests respectively. We also plot the same intervals but for the non-centered versions of these tests. In all cases, an interval covering one indicates statistical significance. The results suggest that the presence of some idiosyncratic terms related to tail risk, such as ReturnSkewCAPM and DownsideBeta, exert additional influence on returns and complement the explanatory power of the five factors.

5.2 Macroeconomics Time Series

In this section, to substantiate Section 5.1, we illustrate our conditional screening test with another empirical application, this time on the macroeconomics dataset FRED-MD introduced in McCracken and Ng (2016) and also later studied in Fan and Gu (2023). The dataset collects d = 127 monthly U.S. macroeconomic variables, such as unemployment rate and real personal income, starting from 1959/01. It is shown in McCracken and Ng (2016) that these variables can be explained well by several latent factors.

Our analysis setup is in general similar to Section C in Fan and Gu (2023). Our target variables are UEMP15T26, TB3SMFFM or TB6SMFFM, and we aim to identify which variables contribute to predicting the target variables beyond the latent factors. The variable UEMP15T26 represents the civilians unemployed for 15 - 26 weeks. The variable TB3SMFFM (TB6SMFFM) measures the 3-month (6-month) treasury bill rate minus the effective federal funds rate. For each target response variable y_{t+1} in {UEMP15T26, TB3SMFFM, TB6SMFFM}, we regress y_{t+1} on $\boldsymbol{x}_t \in \mathbb{R}^{127}$ where \boldsymbol{x}_t is the vector of all variables at the previous month. We choose the n = 330 valid sample pairs { $\boldsymbol{x}_t, y_{t+1}$ } between January 1980 and July 2022. We employ the same neural network fitting parameters as Section 5.1 except that we revert back to a hidden size $k_0 = 16$ as in our simulation



Figure 1: Asset Pricing Dataset: Significance of predictors/idiosyncratic terms at the 5% significance level. Intervals covering one (indicated by the vertical black line) correspond to the significant predictors/idiosyncratic terms. For each variable, six intervals are plotted in the order of: fix-t test, sup test and square test, and within each test the non-centered version is plotted before the centered version.

studies.

For each variable, we plot the same six intervals derived from our tests as in Section 5.1and we recall that an interval covering one indicates statistical significance. Our analysis, based on the variables categorized by McCracken and Ng (2016), underscores the predominant impact of a limited set of variables on the three response variables. As there are 127 variables, it takes 4 figures to display a figure that is similar Figure 1 for each given response variable. Figures F.1 to F.4 in the appendix display the intervals, with the response variable being UEMP15T26 (civilian unemployment). Notably, most variables fail to reach the 5%significance level, except for some variables such as RETAILX, USGOOD, NDMANEMP, and COMPAPFFx. Additional significant variables for the response variables TB3SMFFM and TB6SMFFM are illustrated from Figures F.5 to F.12. As for TB3SMFFM, there are many significant variables. For example, the variables such as IPDMAT in the output and income variable group and DTCTHFNM in the prices group contribute additionally to the factors. TB6SMFFM demonstrates slightly higher susceptibility to variables compared to TB3SMFFM, which, in turn, is notably influenced by variables related to the stock market, employment, interest rates, and money and credit. Our overall findings validate our regression model based on a latent factors plus sparse idiosyncratic terms structure.

6 Conclusion and further work

We have introduced a conditional variable screening test for non-parametric regression using deep neural networks; the inputs to the networks are obtained with the help of a factor model that further enables us to handle very high-dimensional predictors. In our test statistics, we employ high-quality estimators of the partial derivatives of the nonparametric regression function, which could be of independent interest. To demonstrate the versatility of our test, we apply it to assess the adequacy of non-parametric factor regression. An intriguing avenue for further exploration involves extending this framework to dependent data, other statistical machine learning losses, and simultaneous testing for multiple variables. Relevant examples for the latter direction include the ℓ_{∞} statistics proposed by Chen et al. (2022) and the $\ell_2 - \ell_{\infty}$ type statistics discussed by Li et al. (2024).

This paper focuses on the popular feed-forward networks with the ReLU activation function. The derivative irregularity is not limited to the ReLU activation function, and hence, addressing this issue can benefit various neural network classes. Moreover, in our approach, we apply the smoothing procedure after we have obtained a neural network estimator. This is different from other derivative regularizations, for instance, Sobolev training (Czarnecki et al., 2017), that employ roughness penalties during optimization, which could form a potential future topic. Last but not least, we could also potentially conduct different tests for testing higher- or mixed-derivatives. For instance, testing the monotonicity in a variable in the regression is equivalent to testing the sign of the associated partial derivative, a task already studied in the context of an one-dimensional predictor by, for instance, Bowman et al. (1998), Ghosal et al. (2000) and Hall and Huang (2001).

References

- Anthony, M. and Bartlett, P. L. (1999). Neural Network Learning: Theoretical Foundations. Cambridge University Press.
- Baringhaus, L., Ebner, B., and Henze, N. (2017). The limit distribution of weighted l^2 -goodness-of-fit statistics under fixed alternatives, with applications. Annals of the Institute of Statistical Mathematics, 69(5):969–995.
- Barut, E., Fan, J., and Verhasselt, A. (2016). Conditional sure independence screening. *Journal of the American Statistical Association*, 111(515):1266–1277.
- Bowman, A. W., Jones, M. C., and Gijbels, I. (1998). Testing monotonicity of regression. Journal of Computational and Graphical Statistics, 7(4):489–500.
- Candes, E. and Tao, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n. The Annals of Statistics, 35(6):2313-2351.
- Chen, A. Y. and Zimmermann, T. (2022). Open source cross-sectional asset pricing. *Critical Finance Review*, 27(2):207–264.

- Chen, L., Wang, W., and Wu, W. B. (2022). Inference of breakpoints in high-dimensional time series. *Journal of the American Statistical Association*, 117(540):1951–1963.
- Conway, J. B. (1990). A Course in Functional Analysis. Graduate Texts in Mathematics; 96. Springer Science+Business Media, New York, 2nd edition.
- Czarnecki, W. M., Osindero, S., Jaderberg, M., Swirszcz, G., and Pascanu, R. (2017). Sobolev training for neural networks. Advances in Neural Information Processing Systems, 30.
- De Brabanter, K., De Brabanter, J., Gijbels, I., and De Moor, B. (2013). Derivative estimation with local polynomial fitting. *Journal of Machine Learning Research*, 14(1):281–301.
- Dinh, V. and Ho, L. S. T. (2020). Consistent feature selection for analytic deep neural networks. Advances in Neural Information Processing Systems, 33:2420–2431.
- Ebner, B. and Henze, N. (2020). Tests for multivariate normality a critical review with emphasis on weighted l^2 -statistics. *TEST*, 29(4):845–892.
- Elbrächter, D., Perekrestenko, D., Grohs, P., and Bölcskei, H. (2021). Deep neural network approximation theory. *IEEE Transactions on Information Theory*, 67(5):2581–2623.
- Epps, T. W., Singleton, K. J., and Pulley, L. B. (1982). A test of separate families of distributions based on the empirical moment generating function. *Biometrika*, 69(2):391– 399.
- Fan, J., Feng, Y., and Song, R. (2011). Nonparametric independence screening in sparse ultra-high-dimensional additive models. *Journal of the American Statistical Association*, 106(494):544–557.
- Fan, J. and Gijbels, I. (1996). Local Polynomial Modelling and its Applications. Springer.
- Fan, J. and Gu, Y. (2023). Factor augmented sparse throughput deep ReLU neural networks for high dimensional regression. *Journal of the American Statistical Association*, pages 1–15.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360.
- Fan, J., Li, R., Zhang, C.-H., and Zou, H. (2020). Statistical Foundations of Data Science. Chapman and Hall/CRC.
- Fan, J. and Liao, Y. (2022). Learning latent factors from diversified projections and its applications to over-estimated and weak factors. *Journal of the American Statistical* Association, 117(538):909–924.
- Fan, J., Liao, Y., and Yao, J. (2015). Power enhancement in high-dimensional crosssectional tests. *Econometrica*, 83(4):1497–1541.
- Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. Journal of the Royal Statistical Society Series B: Statistical Methodology, 70(5):849–911.

- Fan, J., Zhang, C., and Zhang, J. (2001). Generalized likelihood ratio statistics and wilks phenomenon. *The Annals of statistics*, 29(1):153–193.
- Gasser, T. and Müller, H.-G. (1984). Estimating regression functions and their derivatives by the kernel method. *Scandinavian Journal of Statistics*, pages 171–185.
- Ghosal, S., Sen, A., and van der Vaart, A. W. (2000). Testing monotonicity of regression. *The Annals of Statistics*, pages 1054–1082.
- Gijbels, I. and Goderniaux, A.-C. (2005). Data-driven discontinuity detection in derivatives of a regression function. *Communications in Statistics-Theory and Methods*, 33(4):851– 871.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep Learning. MIT press.
- Hall, P. and Huang, L.-S. (2001). Nonparametric kernel regression subject to monotonicity constraints. *The Annals of Statistics*, 29(3):624–647.
- Ho, L. S. T., Richardson, N., and Tran, G. (2023). Adaptive group Lasso neural network models for functions of few variables and time-dependent data. *Sampling Theory, Signal Processing, and Data Analysis*, 21(2):28.
- Horel, E. and Giesecke, K. (2020). Significance tests for neural networks. Journal of Machine Learning Research, 21(227):1–29.
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., et al. (2023). ChatGPT for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.
- Kennedy, E. H. (2022). Semiparametric doubly robust targeted double machine learning: a review. arXiv preprint arXiv:2203.06469.
- Kohler, M. and Langer, S. (2021). On the rate of convergence of fully connected deep neural network regression estimates. *The Annals of Statistics*, 49(4):2231 2249.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- LeCun, Y., Bengio, Y., and Hinton, G. E. (2015). Deep learning. *Nature*, 521(7553):436–444.
- Li, J., Chen, L., Wang, W., and Wu, W. B. (2024). ℓ^2 inference for change points in highdimensional time series via a two-way MOSUM. The Annals of Statistics, 52(2):602–627.
- Liu, Y. and De Brabanter, K. (2020). Smoothed nonparametric derivative estimation using weighted difference quotients. *Journal of Machine Learning Research*, 21(1):2438–2482.
- McCracken, M. W. and Ng, S. (2016). FRED-MD: A monthly database for macroeconomic research. Journal of Business & Economic Statistics, 34(4):574–589.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Müller, H.-G., Stadtmüller, U., and Schmitt, T. (1987). Bandwidth choice and confidence intervals for derivatives of noisy data. *Biometrika*, 74(4):743–749.
- Quintas-Martinez, V., Chernozhukov, V., Syrgkanis, V., and Newey, W. (2022). RieszNet and ForestRiesz: automatic debiased machine learning with neural nets and random forests. In *International Conference on Machine Learning*, pages 3901–3914. PMLR.
- Rondonotti, V., Marron, J. S., and Park, C. (2007). SiZer for time series: a new approach to the analysis of trends. *Electronic Journal of Statistics*, 1:268–289.
- Schmidt-Hieber, J. (2020). Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4):1875–1897.
- Simonyan, K. and Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.
- Stone, C. J. (1985). Additive regression and other nonparametric models. The Annals of Statistics, 13(2):689–705.
- Telgarsky, M. (2015). Representation benefits of deep feedforward networks. arXiv preprint arXiv:1509.08101.
- Telgarsky, M. (2016). Benefits of depth in neural networks. In Conference on Learning Theory, pages 1517–1539. Proceedings of Machine Learning Research.
- Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., and Ting, D. S. W. (2023). Large language models in medicine. *Nature Medicine*, 29(8):1930–1940.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 58(1):267–288.
- van der Laan, M. J. and Rose, S. (2018). Targeted Learning in Data Science. Springer.
- van der Laan, M. J., Rose, S., et al. (2011). Targeted Learning: Causal Inference for Observational and Experimental Data, volume 4. Springer.
- Yarotsky, D. (2017). Error bounds for approximations with deep ReLU networks. Neural Networks, 94:103–114.
- Zhou, S. and Wolfe, D. A. (2000). On derivative estimation in spline regression. *Statistica Sinica*, pages 93–108.

Supplementary Materials: Conditional nonparametric variable screening test via neural network factor regression

Jianqing Fan Operations Research and Financial Engineering, Princeton University

Weining Wang Department of Economics, Econometrics and Finance, University of Groningen

> Yue Zhao Department of Mathematics, University of York

> > August 17, 2024

Contents

Α	Definitions, miscellaneous results and discussionsA.1DefinitionsA.2Miscellaneous resultsA.3Additional remarksA.4Further motivation on the smoothing operation	3 3 4 6 7
В	Supporting details and preliminary proofs for Section 3B.1Additional technical assumptionsB.2Bias from diversified projection and neural network approximationB.3The stability of conditional expectation in Assumption B.2B.4Useful bounds on key stochastic termsB.5The stochastic term in the proof of Theorem 3.2	 9 11 13 16 19
С	Proofs of main theorems and additional results for Section 3C.1Proof of Theorem 3.1C.2Proof of Theorem 3.2C.3Proof of Proposition 3.3C.4Proof of Theorem 3.4C.5Proof of the size under the null: Theorem 3.5C.5.1Proof for the first term in (C.20)C.5.2Proof for the second term in (C.20)C.5.3Variance estimationC.6Power of the fixed-t test under the alternative	 22 25 27 30 30 31 32 35 37
D	Low-dimensional regime as a special caseD.1 Conditional screening in low-dimensionD.2 Simulation setup for the illustrative example	38 38 39
Е	Figures	39
F,	Additional application results	41

The supplement is organized as follows. Section A collects definitions and some miscellaneous results and discussions. Section B provides supporting details and preliminary proofs for the main theorems in Section 3, while Section C presents the proofs for the said main theorems and the power for the fixed-*t* test. Section D presents the low-dimensional regime as a special case of the high-dimensional regime. Section E collects figures. Section F provides additional graphs for the empirical application in Section 5.2. We introduce the additional shorthand notation that $\langle \cdot, \cdot \rangle$ represents the inner product with respect to the distribution \mathbb{P} (that is, $\langle g_1, g_2 \rangle = \int g_1 g_2 d\mathbb{P}$).

A Definitions, miscellaneous results and discussions

A.1 Definitions

Definition A.1 (Hölder class). Let β and $C_{\mathcal{H}}$ be two positive real numbers and let $[\beta]$ be the largest integer *strictly* less than β . The $(\beta, p, C_{\mathcal{H}})$ -smoothness Hölder class on \mathbb{R}^p is the collection of functions $f : \mathbb{R}^p \to \mathbb{R}$ whose mixed derivatives $(\partial^{\lfloor\beta\rfloor}/\partial z_1^{\omega_1} \dots z_p^{\omega_p})f(\boldsymbol{z})$, where $\boldsymbol{z} = (z_1, \dots, z_p)^{\top}$, $\omega_1, \dots, \omega_p \ge 0$ and $\omega_1 + \dots + \omega_p = \lfloor\beta\rfloor$, all exist, and moreover such derivatives are Hölder-continuous with constant $C_{\mathcal{H}}$ and exponent $\beta - \lfloor\beta\rfloor$, or more precisely, $\left| (\partial^{\lfloor\beta\rfloor}/\partial z_1'^{\omega_1} \dots z_p'^{\omega_p})f(\boldsymbol{z}') - (\partial^{\lfloor\beta\rfloor}/\partial z_1^{\omega_1} \dots z_p'^{\omega_p})f(\boldsymbol{z}) \right| \le C_{\mathcal{H}} \|\boldsymbol{z}' - \boldsymbol{z}\|^{\beta - \lfloor\beta\rfloor}, \forall \boldsymbol{z}, \boldsymbol{z}'.$

Definition A.2 (sub-Gaussian norm; see, for instance, Definition 1 in Adamczak (2008)). Let the function $\psi_2 : [0, \infty) \to [0, \infty)$ be $\psi_2(x) = \exp(x^2) - 1$. Then, the sub-Gaussian norm $\|\cdot\|_{\psi_2}$ of a random variable X is $\|X\|_{\psi_2} = \inf\{\lambda > 0 : \mathbb{E}\psi_2(|X|/\lambda) \leq 1\}.$

Definition A.3 (VC-subgraph class; see, for instance, Section 2.6.2 in van der Vaart and Wellner (1996)). The subgraph of a function $f : \mathcal{X} \to \mathbb{R}$ is the subset of $\mathcal{X} \times \mathbb{R}$ given by $\{(x,t) : x \in \mathcal{X}, t < f(x)\}$. A collection \mathcal{F} of measurable functions is called a VCsubgraph class, or simply a VC-class, if the collection \mathcal{S} of the subgraphs corresponding to all functions $f \in \mathcal{F}$ form a VC class of sets in $\mathcal{X} \times \mathbb{R}$. Then, we define $V_{\mathcal{F}}$ for the function class \mathcal{F} as the VC-index of the associated collection \mathcal{S} of subgraphs. We refer the readers further to, for instance, Section 2.6.1 in van der Vaart and Wellner (1996) for the definitions of VC classes of sets and the accompanying VC-index.

Definition A.4 (VC-type of functions; see, for instance, the beginning of Section 2.1 in Giné and Mason (2007)). Let (S, S) be a measurable space. A class \mathcal{G} of measurable functions on (S, S) is VC-type with respect to an envelope G, if for some $A \ge 3$ and index $v \ge 1$, the covering number $N(\mathcal{G}, L_2(Q), \tau)$ satisfies

$$N(\mathcal{G}, L_2(Q), \tau) \le \left(\frac{A \|G\|_{L_2(Q)}}{\tau}\right)^v, \quad 0 < \tau \le \|G\|_{L_2(Q)}, \tag{A.1}$$

for every probability measure Q on S for which $G \in L_2(Q)$.

Finally we recap a definition coming out from recent deep neural network approximation theory. It is well known that the optimal minimax rate for estimating functions in the $(\beta, \overline{d}, C_{\mathcal{H}})$ -Hölder class (Def. A.1) scales with the sample size n as $n^{-\beta/(2\beta+\overline{d})}$ (Stone, 1982), which is slow when the dimension \overline{d} is large compared to β . To mitigate this issue, Schmidt-Hieber (2020) and Kohler and Langer (2021) showed the excellent approximation power of the neural network class $\mathcal{F}_n(\overline{d})$ in (6) to the *hierarchical composition model* $\mathcal{H}(\overline{d}, l, \mathcal{P}, C_{\mathcal{H}})$ defined in Def. A.5. This model consists of functions that are compositions in $l \in \mathbb{N}$ layers of individual functions in various $(\beta, p, C_{\mathcal{H}})$ -Hölder classes, where the (β, p) tuples belong to a set $\mathcal{P} \subset [1, \infty) \times \mathbb{N}$ (for simplicity we assume a common $C_{\mathcal{H}}$ across the composition functions). Then, under the scaling suggested by Proposition 3.4 in Fan et al. (2022) and recorded in Assumption 3.2, the neural network class $\mathcal{F}(\overline{d})$ can approximate the class $\mathcal{H}(\overline{d}, l, \mathcal{P}, C_{\mathcal{H}})$ adaptively at a rate $(Lk_0)^{-2\kappa}$ up to log factors, where κ is the worst-case dimension-adjusted smoothness measure $\kappa \equiv \min_{(\beta,p)\in\mathcal{P}} \beta/p$. This rate depends only on the *individual* composition functions, and not on the ambient dimension \overline{d} .

Definition A.5. The hierarchical composition model $\mathcal{H}(\overline{d}, l, \mathcal{P}, C_{\mathcal{H}})$ of functions from \mathbb{R}^d to \mathbb{R} , where $\overline{d}, l \in \mathbb{N}$ and $\mathcal{P} \subset [1, \infty) \times \mathbb{N}$, is defined as follows. For l = 1,

$$\mathcal{H}(d, 1, \mathcal{P}, C_{\mathcal{H}}) = \left\{ h : \mathbb{R}^{\overline{d}} \to \mathbb{R} : h(\boldsymbol{x}) = g\left(x_{\pi(1)}, \dots, x_{\pi(p)}\right), \text{ where } g : \mathbb{R}^{p} \to \mathbb{R} \\ \text{ is } (\beta, p, C_{\mathcal{H}}) \text{-H\"older for some } (\beta, p) \in \mathcal{P} \text{ and } C_{\mathcal{H}} < \infty, \text{ and } \pi : \{1, \dots, p\} \to \{1, \dots, \overline{d}\} \right\}.$$

For l > 1, $\mathcal{H}(\overline{d}, l, \mathcal{P}, C_{\mathcal{H}})$ is defined recursively as

$$\mathcal{H}(\overline{d}, l, \mathcal{P}, C_{\mathcal{H}}) = \left\{ h : \mathbb{R}^{d} \to \mathbb{R} : h(\boldsymbol{x}) = g\left(f_{1}(\boldsymbol{x}), \dots, f_{p}(\boldsymbol{x})\right), \text{ where } g : \mathbb{R}^{p} \to \mathbb{R} \\ \text{is } (\beta, p, C_{\mathcal{H}}) \text{-H\"older for some } (\beta, p) \in \mathcal{P} \text{ and } C_{\mathcal{H}} < \infty, \text{ and each } f_{k} \in \mathcal{H}(d, l-1, \mathcal{P}, C_{\mathcal{H}}) \right\}$$

A.2 Miscellaneous results

Lemma A.1. For generic functions $g = g(\widetilde{f}, x_j) : \mathbb{R}^{\overline{r}+1} \to \mathbb{R}$ and $m = m(f, x_j) : \mathbb{R}^{r+1} \to \mathbb{R}$, there exists a constant c^{s} (that does not depend on g or m) such that

$$\int_{\Omega_h} g_j^{\rm s}(\widetilde{\boldsymbol{f}}, x_j)^2 \, \mathrm{d}\mathbb{P} \leqslant \frac{(c^{\rm s})^2}{h^2} \int g(\widetilde{\boldsymbol{f}}, x_j)^2 \mathrm{d}\mathbb{P},$$

$$\int_{\Omega_h} \{g_j^{\rm s}(\widetilde{\boldsymbol{f}}, x_j) - m_j^{\rm s}(\boldsymbol{f}, x_j)\}^2 \, \mathrm{d}\mathbb{P} \leqslant \frac{(c^{\rm s})^2}{h^2} \int \{g(\widetilde{\boldsymbol{f}}, x_j) - m(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P}.$$
(A.2)

Proof. We only prove the first half of (A.2); the second half will follow by a similar argument. We have, starting from (10) and the Cauchy-Schwarz inequality,

$$\begin{split} &\int_{\Omega_h} g_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j)^2 \, \mathrm{d}\mathbb{P} = \int_{\Omega_h} \left\{ \frac{1}{h} \int_{-1}^1 g(\widetilde{\boldsymbol{f}}, x_j - ah) \dot{K}(a) \mathrm{d}a \right\}^2 \, \mathrm{d}\mathbb{P} \\ &\leqslant \frac{1}{h^2} \int_{\Omega_h} \left\{ \int_{-1}^1 g^2(\widetilde{\boldsymbol{f}}, x_j - ah) \mathrm{d}a \right\} \left\{ \int_{-1}^1 \dot{K}(a) \mathrm{d}a \right\} \, \mathrm{d}\mathbb{P} \lesssim \frac{1}{h^2} \int_{\Omega_h} \left\{ \int_{-1}^1 g^2(\widetilde{\boldsymbol{f}}, x_j - ah) \mathrm{d}a \right\} \, \mathrm{d}\mathbb{P} \\ &= \frac{1}{h^2} \int_{-1}^1 \left\{ \int_{\Omega_h} g^2(\widetilde{\boldsymbol{f}}, x_j - ah) \, \mathrm{d}\mathbb{P} \right\} \mathrm{d}a \leqslant \frac{2}{h^2} \|g\|_{L_2}^2. \end{split}$$

Lemma A.2 (Numerical approximation for candidate Riesz representer derivative in Section 2.4). Let N be a large but fixed integer (we used N = 50 throughout our numerical studies) and consider the grid points a_1, \ldots, a_{2N} where $a_l = -1 - \frac{1}{2N} + \frac{l}{N}$ (these grid points are the centers of the 2N equally-sized intervals that form a partition of the interval [-1, 1]). Then, $\alpha_j^s(\tilde{F}_i, X_{i,j})$ in Section 2.4 can be numerically approximated via a middle Riemann sum over the said grid points as

$$\alpha_j^{\mathrm{s}}(\widetilde{F}_i, X_{i,j}) \approx \frac{1}{Nh} \sum_{l=1}^{2N} \dot{K}(a_l) \alpha \left(\widetilde{F}_i, X_{i,j} - a_l h\right).$$

Then, the sample average of the $\alpha_j^{s}(\widetilde{F}_i, X_{i,j})$'s in $\widehat{R}_n^{\text{null}}$ in (12) can be approximated as

$$\frac{1}{n}\sum_{i\in\mathcal{I}_h}\alpha_j^{\mathrm{s}}(\widetilde{F}_i, X_{i,j}) \approx \frac{1}{nNh}\sum_{l=1}^{2N}\dot{K}(a_l)\left\{\sum_{i\in\mathcal{I}_h}\alpha\left(\widetilde{F}_i, X_{i,j}-a_lh\right)\right\}.$$

Proof. We only prove the first equation display, because the second equation display is straightforward given the first. Starting from the right-most term in (10), we have

$$\alpha_j^{\mathrm{s}}(\widetilde{F}_i, X_{i,j}) = \frac{1}{h} \int_{-1}^1 \alpha(\widetilde{F}_i, X_{i,j} - ah) \dot{K}(a) \mathrm{d}a \approx \frac{1}{Nh} \sum_{l=1}^{2N} \dot{K}(a_l) \alpha\left(\widetilde{F}_i, X_{i,j} - a_lh\right),$$

where the approximation follows by replacing the integral with the middle Riemann sum over the grid points in the lemma statement, where each grid point carries a weight 2/(2N) = 1/N.

Lemma A.3 (Existence and analytical form of the Riesz representer in (25)). Under Assumptions B.3 and B.5, the Riesz representer $\alpha_{t,n}^*$ that satisfies (25) exists. Furthermore, if \tilde{F}, X_j admits a joint density $p = p(\tilde{f}, x_j)$, then $\alpha_{t,n}^*$ takes the form

$$\alpha_{t,n}^{*}(\widetilde{\boldsymbol{f}}, x_{j}) = \frac{\int_{\mathcal{B}_{h}} \dot{K}_{h}(z - x_{j}) e^{t \,\Psi\{g_{0,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, z)\}} \dot{\Psi}\{g_{0,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, z)\} p(\widetilde{\boldsymbol{f}}, z) \mathrm{d}z}{p(\widetilde{\boldsymbol{f}}, x_{j})}.$$
(A.3)

Finally, if p is bounded and is differentiable in x_j with the derivative being p_j , then under H_0 and in the limit $h \to 0$, (A.3) simplifies to $-p_j(\tilde{f}, x_j)/p(\tilde{f}, x_j)$.

Proof. We first verify that $\frac{1}{t} \frac{\partial \tilde{\eta}_t^s(g_0)}{\partial g}[v]$ is a bounded linear functional. For any fixed t and h,

$$\frac{\partial \tilde{\eta}_{t}^{s}(g_{0})}{\partial g}[v] = \frac{\partial \tilde{\eta}_{t}^{s}(g_{0} + \tau v)}{\partial \tau}\Big|_{\tau=0} = \left[\frac{\partial}{\partial \tau}\int_{\Omega_{h}} e^{t\Psi\{(g_{0} + \tau v)_{j}^{s}(\tilde{f}, x_{j})\}}d\mathbb{P}\right]\Big|_{\tau=0}$$

$$= \int_{\Omega_{h}} \left[e^{t\Psi\{(g_{0} + \tau v)_{j}^{s}(\tilde{f}, x_{j})\}}\dot{\Psi}\{(g_{0} + \tau v)_{j}^{s}(\tilde{f}, x_{j})\}\right]\Big|_{\tau=0} tv_{j}^{s}(\tilde{f}, x_{j})d\mathbb{P}$$

$$= t\int_{\Omega_{h}} e^{t\Psi\{g_{0,j}^{s}(\tilde{f}, x_{j})\}}\dot{\Psi}\{g_{0,j}^{s}(\tilde{f}, x_{j})\}v_{j}^{s}(\tilde{f}, x_{j})d\mathbb{P}.$$
(A.4)

Thus, the functional $\frac{1}{t} \frac{\partial \tilde{\eta}_t^{s}(g_0)}{\partial g}[v]$ is clearly linear in its input and, by the boundedness of Ψ , $\dot{\Psi}$ and Lemma A.1, is also bounded. Therefore, the Riesz representer that satisfies (25) exists (Conway, 1990, Theorem 3.4, Chapter 1). As a side remark, the modified functional that, when acting on any v, always returns $\int_{\Omega_h} v_j^{s}(\tilde{f}, x_j) d\mathbb{P}$ (the last term in (25)) is also a bounded linear functional. Hence, this modified functional admits its own Riesz representer. It's also easy to show that the latter Riesz representer minimizes the population loss $R_n^{\text{null}} \equiv \int \hat{R}_n^{\text{null}}(\alpha) d\mathbb{P}$. Thus the latter Riesz representer is precisely α_n^{null} which we have defined in Section 3.2.1 as the population limit of $\hat{\alpha}_n$.

Next, starting from the first two steps of (25), with the assumed density p,

$$\begin{split} t \langle v, \alpha_{t,n}^* \rangle &= \frac{\partial \tilde{\eta}_{t}^{\mathrm{s}}(g_{0})}{\partial g} [v] = t \int_{[-b,b]^{\overline{r}}} \int_{\mathcal{B}_{h}} \{ \int_{[-b,b]} v(\tilde{f},z) \dot{K}_{h}(x_{j}-z) \mathrm{d}z \} e^{t \, \Psi\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\}} \dot{\Psi}\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\} p(\tilde{f},x_{j}) \mathrm{d}x_{j} \mathrm{d}\tilde{f} \\ &= t \int_{[-b,b]^{\overline{r}+1}} v(\tilde{f},z) \{ \int_{\mathcal{B}_{h}} \dot{K}_{h}(x_{j}-z) e^{t \, \Psi\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\}} \dot{\Psi}\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\} p(\tilde{f},x_{j}) \mathrm{d}x_{j} \mathrm{d}z \mathrm{d}\tilde{f} \\ &= t \int_{[-b,b]^{\overline{r}+1}} v(\tilde{f},z) \left\{ \frac{\int_{\mathcal{B}_{h}} \dot{K}_{h}(x_{j}-z) e^{t \, \Psi\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\}} \dot{\Psi}\{g_{0,j}^{\mathrm{s}}(\tilde{f},x_{j})\} p(\tilde{f},x_{j}) \mathrm{d}x_{j} }{p(\tilde{f},z)} \right\} p(\tilde{f},z) \mathrm{d}z \mathrm{d}\tilde{f}. \end{split}$$

After interchanging the dummy integration variables z and x_j in the last line above, the term in the curly bracket then gives the form of $\alpha_{t,n}^*$ in (A.3). Finally,

$$\alpha_{t,n}^{*}(\widetilde{\boldsymbol{f}},x_{j}) = \frac{\int_{-b+h}^{b-h} \dot{K}_{h}\left(z-x_{j}\right) e^{t\Psi\{g_{0,j}^{s}(\widetilde{\boldsymbol{f}},z)\}} \dot{\Psi}\{g_{0,j}^{s}(\widetilde{\boldsymbol{f}},z)\} p(\widetilde{\boldsymbol{f}},z) \mathrm{d}z}{p(\widetilde{\boldsymbol{f}},z)} \stackrel{\mathrm{under}}{=} {}_{H_{0}} \frac{\int_{-b+h}^{b-h} \dot{K}_{h}\left(z-x_{j}\right) p(\widetilde{\boldsymbol{f}},z) \mathrm{d}z}{p(\widetilde{\boldsymbol{f}},z)} \mathrm{d}z}$$

For $x_j \in [-b + 2h, b - 2h]$, the numerator can be treated as

$$\begin{split} \int_{-b+h}^{b-h} \dot{K}_h \left(z - x_j \right) p(\widetilde{\boldsymbol{f}}, z) \mathrm{d}z &= -\int_{x_j-h}^{x_j+h} \dot{K}_h \left(x_j - z \right) p(\widetilde{\boldsymbol{f}}, z) \mathrm{d}z = -\frac{\partial}{\partial x_j} \int_{x_j-h}^{x_j+h} K_h \left(x_j - z \right) p(\widetilde{\boldsymbol{f}}, z) \mathrm{d}z \\ &= -\frac{\partial}{\partial x_j} \int_{x_j-h}^{x_j+h} K_h \left(z \right) p(\widetilde{\boldsymbol{f}}, x_j - z) \mathrm{d}z = -\int_{x_j-h}^{x_j+h} K_h \left(z \right) p_j(\widetilde{\boldsymbol{f}}, x_j - z) \mathrm{d}z \xrightarrow{h \to 0} -p_j(\widetilde{\boldsymbol{f}}, x_j), \end{split}$$

where we have applied twice differentiation under the integral sign (Protter and Morrey, Jr., 2012, Theorem 3, Chapter 8), followed by the usual approximation to the identity theory (Stein and Shakarchi, 2009, Theorem 2.1). This completes the proof of the Lemma. \Box

A.3 Additional remarks

Remark A.1 (Specification test in factor models). As discussed in Section 2.1, our hypotheses (3) and (4) are equivalent to determining whether the *j*-th idiosyncratic term μ_j of **u** (defined in (1)) contributes to the regression function at the population level. In the framework of Fan and Liao (2022), this test is further similar to testing their Factor Augmented Regression (FAR) model against their Factor Augmented Sparse Throughput (FAST) model for high-dimensional regression. To establish the aforementioned equivalence, let's define the regression function

$$m_0^{\dagger}(\boldsymbol{F},\mu_j) = \mathbb{E}[Y|\boldsymbol{F},\mu_j]$$

and the associated partial derivative with respect to the last argument:

$$m_{0,j}^{\dagger}(\boldsymbol{f}, u_j) = \frac{\partial}{\partial u_j} m_0^{\dagger}(\boldsymbol{f}, u_j).$$
(A.5)

Then, similar to our hypotheses (3) and (4), the new hypotheses for μ_j can be formulated as

$$H_0^{\dagger}: m_{0,j}^{\dagger}(\boldsymbol{f}, u_j) = 0 \text{ for all } \boldsymbol{f}, u_j, \quad \text{against} \quad H_A^{\dagger}: m_{0,j}^{\dagger}(\boldsymbol{f}, u_j) \neq 0 \text{ for some } \boldsymbol{f}, u_j.$$
(A.6)

Let \mathbf{B}_{j} . denote the *j*-th row of the loading matrix \mathbf{B} , so $X_j = \mathbf{B}_{j} \cdot \mathbf{F} + \mu_j$. Clearly, \mathbf{F}, X_j and \mathbf{F}, μ_j generate the same σ -algebra, and hence the functions m_0 and m_0^{\dagger} as conditional expectations are essentially a re-writing of each other with different arguments. Specifically,

$$m_0(\boldsymbol{f}, x_j) = m_0^{\dagger}(\boldsymbol{f}, x_j - \mathbf{B}_j, \boldsymbol{f}) = m_0^{\dagger}(\boldsymbol{f}, u_j)$$

Thus, the partial derivatives $m_{0,j}$ and $m_{0,j}^{\dagger}$ are linked as follows: at $x_j - \mathbf{B}_j \mathbf{f} = u_j$,

$$m_{0,j}(\boldsymbol{f}, x_j) = \frac{\partial}{\partial x_j} m_0(\boldsymbol{f}, x_j) = \frac{\partial}{\partial x_j} m_0^{\dagger}(\boldsymbol{f}, x_j - \mathbf{B}_j \cdot \boldsymbol{f}) = m_{0,j}^{\dagger}(\boldsymbol{f}, u_j).$$

Therefore, at least at the population level, our conditional screening hypotheses (3) and (4) for $m_{0,j}$ are equivalent to the similar ones in (A.6) for $m_{0,j}^{\dagger}$.

Remark A.2 (Relationship between (2) and the full model). Note that our reduced regression model (2) obviously does not preclude the "full" regression model $Y = m_0^*(\mathbf{F}, \mathbf{X}) + \epsilon^*$ where $m_0^*(\mathbf{F}, \mathbf{X}) = \mathbb{E}[Y|\mathbf{F}, \mathbf{X}]$ is the full conditional expectation. It is apparent that our reduced model and the full model are linked as

$$\epsilon = \epsilon^* + \mathbb{E}[Y|\boldsymbol{F}, \boldsymbol{X}] - \mathbb{E}[Y|\boldsymbol{F}, X_j] = \epsilon^* + m_0^*(\boldsymbol{F}, \boldsymbol{X}) - m_0(\boldsymbol{F}, X_j).$$

Moreover, by the definition of the conditional expectation, ϵ is centered and uncorrelated with any function of \mathbf{F} and X_i .

Remark A.3 (Boundary region). Technically, we need to exclude the boundary region, which under Assumption 3.1 is the union $[-b, -b + h] \cup [b - h, b]$, from our smoothing operation, because smoothing at a x_j in this region will extend beyond the support of X_j which is problematic. Thus the sets \mathcal{B}_h , \mathcal{I}_h , Ω_h relevant to the interior regions in different contexts were introduced. For instance, the quantity $\int_{\Omega_h} g d\mathbb{P}$ signifies the integral of the function g in the interior region only.

Remark A.4 (Cross validation for bandwidth selection). In this remark we provide some brief guideline on how to choose an optimal bandwidth h through cross validation. Suppose that we split the data evenly into K folds (for instance, K = 5) indexed by k = 1, ..., K, and denote the collection of sample indices $i \in \{1, ..., n\}$ within the k-th fold by \mathcal{I}_k . We let $\hat{g}_n^{(-k)} = \hat{g}_n^{(-k)}(\tilde{f}, x_j)$ be the regression function estimator similar to that in (7), except that here we exclude the samples in the k-th fold; then, let $\hat{g}_{n,j,h}^{(-k),s} = \hat{g}_{n,j,h}^{(-k),s}(\tilde{f}, x_j)$ be the corresponding estimator for the partial derivative (with respect to x_j) that is smoothed at bandwidth h. Next, let $\hat{g}_n^{(k)}$ be the regression function estimator using only the sample in the k-th fold, and let $\hat{g}_{n,j}^{(k)}$ be the corresponding unsmoothed estimator for the partial derivative (with respect to x_j). Then, we define the cross validation error at h as CV(h) = $\sum_{k=1}^{K} \sum_{i \in \mathcal{I}_k} [\hat{g}_{n,j,h}^{(-k),s}(\tilde{F}_i, X_{i,j}) - \hat{g}_{n,j}^{(k)}(\tilde{F}_i, X_{i,j})]^2$. Finally, we could look for the value of h that minimizes CV(h) as the optimal h.

A.4 Further motivation on the smoothing operation

In this section, we first provide detailed descriptions for our observation in Section 2.3 on the sawtooth function ζ_L which demonstrates the intricacy of derivative estimation with deep neural networks, and then further motivate our smoothing operation. By Proposition 2 in Yarotsky (2017), the triangular function $\zeta : [0, 1] \rightarrow [0, 1]$ defined as

$$\zeta(t) = \begin{cases} 2t, & 0 \le t < 1/2, \\ 2(1-t), & 1/2 \le t \le 1, \\ 0 & \text{otherwise} \end{cases}$$

is given in terms of the ReLU function σ by $\zeta(t) = 2\sigma(t) - 4\sigma(t-1/2) + 2\sigma(t-1)$. Next, let $\zeta_L(\cdot) = \frac{2}{2^L}\zeta_L \circ \zeta_L \circ \cdots \circ \zeta_L$ denote the *L*-fold iteration of ζ that is further scaled by an overall factor $\frac{2}{2^L}$. The function ζ_L , described in Lemma 2.4 in Telgarsky (2015), is often referred to as the "sawtooth" function. (To be precise, in that reference, ζ_L is defined without the scaling factor $\frac{2}{2^L}$.) The function ζ_L is visually represented in Figure E.1 that in particular illustrates the increasingly "oscillatory" behavior of ζ_L for increasing values of *L*. By Proposition 2 in Yarotsky (2017), it's feasible to implement ζ_L with a neural network of no more than L hidden layers and three nodes per layer. (See Fig. 2(c) in Yarotsky (2017) for a related construction.)

The function ζ_L essentially consists of 2^L piecewise linear segments with each segment having a slope of constant magnitude 2. Consequently, the L_2 norm of ζ_L over the interval [0, 1] with respect to the Lebesgue measure is $\{2^L \int_0^{1/(2^L)} (2t)^2 dt\}^{1/2} \sim 1/2^L$, which vanishes as L tends to infinity. Meanwhile, let the first-order derivative of ζ_L be $\dot{\zeta}_L$. Because the magnitude of $\dot{\zeta}_L$ is always 2, its L_2 norm (with respect to the same Lebesgue measure on [0, 1]) is also 2. Thus, the L_2 norm of the derivative function $\dot{\zeta}_L$ remains a constant, even though the L_2 norm of the function ζ_L itself converges exponentially in L to zero, yielding the observation in Section 2.3.

Clearly, this example implies that the L_2 convergence between a neural network function and a target, for instance from \hat{g}_n to m_0 , does not necessarily imply the convergence at the same rate between their partial derivatives, for instance from $\hat{g}_{n,j}$ to $m_{0,j}$. Moreover, this discrepancy becomes more apparent as the depth L increases, but this is precisely the regime of interest for deep neural networks. Mathematically, this phenomenon can be understood as the absence of a general reverse Poincaré inequality. The inequality itself conventionally bounds the function norm by the norm of the derivative of the said function (see Theorem 13.27 in Leoni (2017) for an exact formulation), but the reverse direction does not always hold, as demonstrated by our simple example.

It's important to clarify that we're not asserting that *no* neural network can approximate the first-order derivatives of smooth functions (Gühring et al., 2020). In our example above, we could easily select a neural network that consistently outputs zero, thereby perfectly approximating both the identically zero function and its derivative. However, in practical scenarios involving observations with noise, there's no guarantee that such a derivative-approximating neural network will be chosen. Furthermore, as previously noted, for approximating higher-order derivatives, any ReLU-based neural network will completely fail if we simply settle for the derivatives of the same order (which will always be zero almost everywhere) on the network function directly.

Having demonstrated the irregularities that the derivatives of neural networks may exhibit, we now motivate our smoothing operation. Our first motivation is the convergence performance obtained from the smoothing operation and the generality of the operation. As mentioned in Section 1.2, our smoothing operation proposed in (10) is applicable to general machine learning estimators. Specifically, our Lemma A.4 shows that, if we start from a generic regression estimator \hat{q} for a generic regression function target m, then by using the simple smoothing operation in (10), we can easily recover a convergence rate from the *derivative* estimator \hat{g}_i^s to the target derivative m_j through the convergence rate from \hat{g}_n to m_0 , up to a factor 1/h and a bias term; thus the two rates can be made comparable as long as the bandwidth h does not approach zero too fast. Therefore, we have translated the power of neural networks for regression function estimation into an analogous result for derivative estimation through smoothing. The bias term occurs naturally because \hat{g}_{i}^{s} is the smoothed version of the naive derivative estimator \hat{g}_i , and hence its target is the smoothed target derivative m_j^s instead of the non-smoothed target derivative m_j . Later, we will apply our lemma on our deep neural network estimators, which will replace \hat{g} here with \hat{g}_n in (7) or \check{g}_n in (14).

Lemma A.4. Suppose that Assumptions 3.1(*i*) and B.3 hold. Let $\hat{g} = \hat{g}(\tilde{f}, x_j) : \mathbb{R}^{\bar{r}+1} \to \mathbb{R}$ be a generic estimator that estimates a generic target $m = m(f, x_j) : \mathbb{R}^{r+1} \to \mathbb{R}$ with a L_2 rate of r_n on an event \mathcal{A} , that is

$$\{ \int (\hat{g} - m)^2 \, \mathrm{d}\mathbb{P} \}^{1/2} \leqslant r_n \text{ on the event } \mathcal{A}.$$
(A.7)

Moreover, let m_j^s and m_j be respectively the smoothed and non-smoothed derivatives (with respect to the last argument) of the target function m, and suppose that m_j^s exhibits a bias with respect to m_j as follows:

$$\{\int_{\Omega_h} (m_j^{\mathrm{s}} - m_j)^2 \,\mathrm{d}\mathbb{P}\}^{1/2} \leqslant Bias_n.$$
(A.8)

Then, there exists a constant c^{s} (that does not depend on g or m) such that the smoothed derivative estimator \hat{g}_{i}^{s} estimates the non-smoothed m_{i} with the following rate:

$$\{\int_{\Omega_h} (\widehat{g}_j^s - m_j)^2 \, \mathrm{d}\mathbb{P}\}^{1/2} \leqslant \frac{c^s}{h} r_n + Bias_n \text{ on the event } \mathcal{A}.$$
 (A.9)

Proof. By the triangle inequality, the left-hand side of (A.9) can be bounded as

$$\{\int_{\Omega_h} (\hat{g}_j^{\rm s} - m_j)^2 \,\mathrm{d}\mathbb{P}\}^{1/2} \leqslant \{\int_{\Omega_h} (\hat{g}_j^{\rm s} - m_j^{\rm s})^2 \,\mathrm{d}\mathbb{P}\}^{1/2} + \{\int_{\Omega_h} (\hat{m}_j^{\rm s} - m_j)^2 \,\mathrm{d}\mathbb{P}\}^{1/2}$$

The first term on the right-hand side above can be bounded by Lemma A.1 in conjunction with the rate in (A.7), and the second term on the right-hand side above is simply bounded by the condition given in (A.8). \Box

The second motivation for our smoothing operation comes from traditional non-parametric regression. It is well known that in the latter setting, the optimal bandwidth for the derivative estimator of either the density or regression function should be larger than the optimal bandwidth for the estimator of the density or regression function itself (see, for instance, (Gasser and Müller, 1984, Theorem 1), (Hall and Marron, 1987, Theorem 3.2) and (Härdle et al., 1990, p. 226)), due to the former estimator possessing a larger variance at the same bandwidth. Another example in non-parametric regression calling for over-smoothing comes from the so-called irrelevant regressors (Li and Racine, 2006, Section 2.2.4), which are those predictors that do not materially contribute to the regression outcome. Irrelevant regressors are not about derivative estimation per se, but by their very definition they are similar to those coordinates X_k that satisfy the null hypothesis (3) in our screening test. In kernel regression, it is known that the optimal bandwidth for the irrelevant regressors should diverge to infinity in order to "weed out" these regressors more efficiently. (Furthermore, if a regressor is not irrelevant, but contributes linearly to the regression outcome, it's conjectured that its associated optimal bandwidth should diverge to infinity as well (Li and Racine, 2006, Section 2.4.1).) Thus, from the point of view of both derivative estimation and screening in traditional non-parametric regression, it is not surprising that smoothing could benefit our derivative estimation and screening test.

B Supporting details and preliminary proofs for Section **3**

B.1 Additional technical assumptions

Assumption B.1 (Additional smoothness of m_0 and conditions on noise). The true regression function $m_0(\mathbf{f}, x_i)$ is Lipschitz in the first argument, precisely in the sense that $|m_0(\mathbf{f}, x_i) - \mathbf{f}_i|$

 $m_0(\mathbf{f}', x_j)| \leq C_L \|\mathbf{f} - \mathbf{f}'\|_2$ for a common positive constant C_L for all x_j , \mathbf{f} and \mathbf{f}' ; in addition, for a positive constant M'_{∞} , the derivative of m_0 in the last argument, namely $m_{0,j}$, exists everywhere and satisfies $\sup_{(\mathbf{f}, x_j) \in [-b,b]^{r+1}} |m_{0,j}(\mathbf{f}, x_j)| < M'_{\infty}$. The noise $\epsilon = Y - m_0(\mathbf{F}, X_j)$ is sub-Gaussian (see Def. A.2) with sub-Gaussian norm $C_{\epsilon} < \infty$.

Assumption B.2 (Diversified projection matrix). The diversified projection matrix $\mathbf{W} = (\mathbf{W}_{j,k})_{1 \leq j \leq d, 1 \leq k \leq \overline{r}} \in \mathbb{R}^{d \times \overline{r}}$ is exogenous, so is independent of $(Y_i, \mathbf{F}_i, \mathbf{X}_i)$, $i = 1, \ldots, n$. For some constants $c_{\mathbf{W},1}$, $c_{\mathbf{W},2}$, and $c_{\mathbf{W},3}$, $\max_{j,k} |\mathbf{W}_{j,k}| \leq c_{\mathbf{W},1}$, and the minimum and maximum singular values of \mathbf{H} (recall that $\mathbf{H} = d^{-1}\mathbf{W}^{\top}\mathbf{B}$), denoted by $v_{\min}(\mathbf{H})$ and $v_{\max}(\mathbf{H})$ respectively, satisfy $0 < c_{\mathbf{W},2} \leq v_{\min}(\mathbf{H}) \leq v_{\max}(\mathbf{H}) \leq c_{\mathbf{W},3}$. Finally,

$$\mathbb{E}[\{\mathbb{E}[Y|\widetilde{\boldsymbol{F}}, X_j] - m_0(\boldsymbol{F}, X_j)\}^2] \lesssim \delta_{\rm f}^2.$$
(B.1)

Assumption B.2 above imposes conditions on W beyond exogeneity to ensure in particular $\mathbf{H}^{\dagger} \widetilde{\mathbf{F}} \approx \mathbf{F}$. Moreover, Eq. (B.1) in Assumption B.2 is a high-level condition regarding the two different conditional expectations $\mathbb{E}[Y|\widetilde{\mathbf{F}}, X_j]$ and (from (2)) $m_0(\mathbf{F}, X_j) = \mathbb{E}[Y|\mathbf{F}, X_j]$. This condition is necessary because while our regression model (2) involves m_0 , our estimation procedures instead rely on the (sample of the) diversified factors $\widetilde{\mathbf{F}}$ and X_j , so $\mathbb{E}[Y|\widetilde{\mathbf{F}}, X_j]$ becomes a more appropriate measure of the "center" of Y. Intuitively, because \mathbf{F}, X_j and $\widetilde{\mathbf{F}}, X_j$ should generate similar σ -algebras, $\mathbb{E}[Y|\widetilde{\mathbf{F}}, X_j]$ should be close to $m_0(\mathbf{F}, X_j)$ as characterized by (B.1) (Bryc and Smolenski, 1992). We will provide conditions in Lemma B.5 for the validity of (B.1).

Assumption B.3 (Smoothing kernel). The kernel $K : \mathbb{R} \to \mathbb{R}$ is compactly supported on [-1,1], satisfies $\int_{-1}^{1} K(a) da = 1$ and vanishes at the boundary points -1 and 1, and is continuously differentiable with uniformly bounded derivative \dot{K} . The bandwidth h satisfies $h \leq b$ (and thus the interior sets \mathcal{B}_h , \mathcal{I}_h and Ω_h defined in Section 2.3 are not degenerate).

Assumption B.4 (Miscellaneous conditions for estimating $\alpha_{t,n}^*$). The class of functions $\{\alpha_{t,n}^* : t \in \mathcal{T}_{\delta}\}$ indexed by t is of VC-type (see Definition A.4) with an index of at most $V_{\mathcal{F}_n(\bar{r}+1)}$, and is uniformly bounded in magnitude by $M_{\infty} < \infty$. Moreover, for each $k \in \{1, \ldots, \bar{r}\}$, for \mathbf{H}_k the k-th row of \mathbf{H} , $\mathbb{P}(|\mathbf{H}_k, \mathbf{F}| > c_b b/2) \leq \delta_{\mathrm{f}}^2/\bar{r}$.

In Assumption B.4, $V_{\mathcal{F}_n(\bar{r}+1)}$ is the VC-index of the VC-subgraph class of functions $\mathcal{F}_n(\bar{r}+1)$ defined in terms of subgraphs as in Definition A.3. On the other hand, the VC-type assumption on the class $\{\alpha_{t,n}^* : t \in \mathcal{T}_\delta\}$ refers to Definition A.4; the "index" therein is not explicitly defined in terms of subgraphs and Assumption B.4 simply states that this index is the number $V_{\mathcal{F}_n(\bar{r}+1)}$. Technically, the VC-subgraph and the VC-type classes of functions are related: a VC-subgraph class of functions with a VC-index v admits a similar covering number as a VC-type class of functions that admits the same index v. However, conceptually, these two classes of functions are not equal. It's also worth noting that in Assumption B.4, $\|\mathbf{H}_k\|_2$ is upper bounded by $v_{\max}(\mathbf{H})$ which is in turn bounded by a constant as stipulated in Assumption B.2. Thus, the tail condition on $\mathbf{H}_k.\mathbf{F}$ in Assumption B.4 is satisfied if \mathbf{F} is coordinate-wise bounded.

Assumption B.5 (Truncation function). The truncated function Ψ is twice differentiable, bounded, and have bounded first- and second-order derivatives $\dot{\Psi}$ and $\ddot{\Psi}$ respectively. In addition, Ψ satisfies $\Psi(t) = t$ for $t \in [-\|K\|_{L_1}M'_{\infty}, \|K\|_{L_1}M'_{\infty}]$ (where M'_{∞} is defined in Assumption B.1 and $\|K\|_{L_1}$ is the L_1 norm of the kernel K). For an example of a truncation function Ψ satisfying Assumption B.5, we can simply "insert" a linear segment of length 2γ , where $\gamma > ||K||_{L_1}M'_{\infty}$, around the origin into a centered and scaled logistic function; specifically, this Ψ is given by $\Psi(t) = t$ for $|t| \leq \gamma$ and $\Psi(t) = \operatorname{sgn}(t)\gamma - 1/2 + 1/[1 + \exp\{-4(t - \operatorname{sgn}(t)\gamma)\}]$ for $|t| > \gamma$.

B.2 Bias from diversified projection and neural network approximation

$$\widetilde{m}_n = \operatorname{argmin}_{m=m(\boldsymbol{f}, x_j) \in \mathcal{F}_n(r+1)} \sup_{(\boldsymbol{f}, x_j) \in [-2b, 2b]^r \times [-b, b]} |(m - m_0)(\boldsymbol{f}, x_j)|, \quad (B.2)$$

and second, in Lemma B.4 we show that analogously the true Riesz representer under the null hypothesis α_n^{null} is well approximated by

$$\widetilde{\alpha}_n = \operatorname{argmin}_{\alpha \in \mathcal{F}_n(\overline{r}+1)} \sup_{(\widetilde{f}, x_j) \in [-c_b b, c_b b]^{\overline{r}} \times [-b, b]} |(\alpha - \alpha_n^{\operatorname{null}})(\widetilde{f}, x_j)|.$$
(B.3)

We start with Lemma B.1. Let $\nu_{a,n}$ be the approximation rate in the sup norm attained by \widetilde{m}_n in (B.2), that is

$$\nu_{\mathbf{a},n} = \sup_{(\boldsymbol{f},x_j) \in [-2b,2b]^r \times [-b,b]} |(\widetilde{m}_n - m_0)(\boldsymbol{f},x_j)|.$$

As can be seen from the lemma and its proof below, the price we pay for the approximation is decomposed into the term $\delta_{\rm f}$, which in fact diminishes with respect to the ambient dimension d and which results from approximating m_0 using diversified factors through $g_0 = g_0(\tilde{f}, x_j)$ given in (8) (Lemma B.2), and the term $\nu_{\rm a,n}$, which results from subsequently approximating g_0 by the neural network function \tilde{g}_n (Lemma B.3). Our lemma is a straightforward extension of Ineq. (F.6) in the appendix of Fan and Gu (2023) by including the variable of interest X_j . For completeness, we will provide a sketch of the proof.

Lemma B.1 (Modification of Inequality (F.6) in Fan and Gu (2023)). Under Assumptions 3.1 to 3.3 and B.1 to B.2,

$$\left[\int \{ \widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, x_j) \}^2 \mathrm{d}\mathbb{P} \right]^{1/2} \lesssim \delta_{\mathrm{f}} + \nu_{\mathrm{a}, n}.$$

Proof. First, we have the easy bound

$$\begin{split} & [\int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P}]^{1/2} \\ & \leq [\int \{g_0(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P}]^{1/2} + [\int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 \mathrm{d}\mathbb{P}]^{1/2} \end{split}$$

The first and second term in the last line above are in turn handled by Lemma B.2 and Lemma B.3 below respectively, whose rates when combined yield the conclusion of the lemma. $\hfill \Box$

Lemma B.2. Under the same assumptions as Lemma B.1,

$$\left[\int \{g_0(\widetilde{\boldsymbol{f}}, x_j) - m_0(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P} \right]^{1/2} \lesssim \delta_{\mathrm{f}}.$$

Proof. By the definition of diversified projection matrix, $\operatorname{rank}(\mathbf{H}_{j}^{\dagger}) = r$, which implies $\mathbf{H}^{\dagger}\mathbf{H} = \mathbf{I}_{r}$. Starting from the decomposition

$$\widetilde{\boldsymbol{F}}_{i} = d^{-1} \mathbf{W}^{\top} \boldsymbol{X}_{i} = d^{-1} \mathbf{W}^{\top} \mathbf{B} \boldsymbol{F}_{i} + d^{-1} \mathbf{W}^{\top} \boldsymbol{u}_{i} = \mathbf{H} \boldsymbol{F}_{i} + d^{-1} \mathbf{W}^{\top} \boldsymbol{u}_{i}, \qquad (B.4)$$

we have

$$m_0(\mathbf{F}_i, X_{i,j}) - g_0(\widetilde{\mathbf{F}}_i, X_{i,j})$$

$$= m_0(\mathbf{F}_i, X_{i,j}) - g_0(\mathbf{H}\mathbf{F}_i + d^{-1}\mathbf{W}^{\top}\mathbf{u}_i, X_{i,j})$$

$$= m_0(\mathbf{F}_i, X_{i,j}) - m_0(\mathbf{H}^{\dagger}\mathbf{H}\mathbf{F}_i + d^{-1}\mathbf{H}^{\dagger}\mathbf{W}^{\top}\mathbf{u}_i, X_{i,j}) \quad \text{by the definition of } g_0 \text{ in } (8)$$

$$= m_0(\mathbf{F}_i, X_{i,j}) - m_0(\mathbf{F}_i + d^{-1}\mathbf{H}^{\dagger}\mathbf{W}^{\top}\mathbf{u}_i, X_{i,j}) \quad \text{by the definition of } \mathbf{H}^{\dagger}. \tag{B.5}$$

By Assumption B.1, $m_0(\boldsymbol{f}, x_j)$ is Lipschitz in the sense that $|m_0(\boldsymbol{f}, x_j) - m_0(\boldsymbol{f}', x_j)| \leq C_L \|\boldsymbol{f} - \boldsymbol{f}'\|$ (for all x_j). It follows from this and (B.5) that

$$|m_0(\boldsymbol{F}_i, X_{i,j}) - g_0(\widetilde{\boldsymbol{F}}_i, X_{i,j})| \leq C_{\mathrm{L}} d^{-1} \left\| \mathbf{H}^{\dagger} \mathbf{W}^{\top} \boldsymbol{u}_i \right\|_2 \leq C_{\mathrm{L}} d^{-1} \left\| \mathbf{H}^{\dagger} \right\|_{\mathrm{op}} \left\| \mathbf{W}^{\top} \boldsymbol{u}_i \right\|_2.$$
(B.6)

By the derivation of (F.15) in Fan and Gu (2023) and Assumption B.2,

$$\|\mathbf{H}^{\dagger}\|_{\text{op}} \leqslant \left\{ v_{\min}(\mathbf{H}^{\dagger}) \right\}^{-1} \lesssim 1.$$
(B.7)

Next we bound $||d^{-1}\mathbf{W}^{\top}\boldsymbol{u}_i||_2$ in (B.6). It follows from the linearity of expectation and then Assumption 3.3 that

$$\mathbb{E}[\|\mathbf{W}^{\top}\boldsymbol{u}_{i}\|^{2}] = \mathbb{E}\left[\sum_{k=1}^{\overline{r}} (\sum_{j=1}^{d} W_{j,k}\mu_{i,j})^{2}\right] = \sum_{k=1}^{\overline{r}} \sum_{j=1}^{d} W_{j,k}^{2} \mathbb{E}\left[\mu_{i,j}^{2}\right] + \sum_{j\neq j'} W_{j,k} W_{j',k} \mathbb{E}\left[\mu_{i,j}\mu_{i,j'}\right]$$
$$\leqslant \overline{r} \max_{j,k} |W_{j,k}| \sum_{j=1}^{d} \mathbb{E}\left[\mu_{i,j}^{2}\right] + \overline{r} \max_{j,k} |W_{j,k}|^{2} \sum_{j\neq j'} |\mathbb{E}\left[\mu_{i,j}\mu_{i,j'}\right]| \lesssim \overline{r} d_{\mu}.$$
(B.8)

Finally, we conclude by taking the expectation of the square of both sides of (B.6) and subsequently applying (B.7) and (B.8).

Lemma B.3. Under the same assumptions as Lemma B.1,

$$\left[\int \{ \widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j) \}^2 \mathrm{d}\mathbb{P} \right]^{1/2} \lesssim \nu_{\mathrm{a}, n} + \delta_{\mathrm{f}}.$$

Proof. Following essentially the same derivation on p. 10 in the appendix of Fan and Gu (2023),

$$\begin{split} & \int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 d\mathbb{P} = \int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 \mathbb{1}\{\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}} \in [-2b, 2b]^r\} d\mathbb{P} \\ & + \int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 \mathbb{1}\{\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}} \notin [-2b, 2b]^r\} d\mathbb{P}. \end{split}$$

By the definition of \tilde{g}_n , g_0 and $\nu_{a,n}$, the first term on the right-hand side above is reduced to

$$\begin{split} & \int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 \mathbb{1}\{\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}} \in [-2b, 2b]^r\} d\mathbb{P} \\ &= \int \{\widetilde{m}_n(\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}}, x_j) - m_0(\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}}, x_j)\}^2 \mathbb{1}\{\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}} \in [-2b, 2b]^r\} d\mathbb{P} \leqslant \nu_{\mathbf{a}, n}^2 \end{split}$$

Meanwhile, for the second term, by Assumption B.1, Markov's inequality, and finally the reasoning at the top of p. 11 in the appendix of Fan and Gu (2023),

$$\begin{split} & \int \{\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j) - g_0(\widetilde{\boldsymbol{f}}, x_j)\}^2 \mathbb{1}\{\mathbf{H}^{\dagger} \widetilde{\boldsymbol{f}} \notin [-2b, 2b]^r\} d\mathbb{P} \\ & \leq (\|\widetilde{g}_n(\widetilde{\boldsymbol{f}}, x_j)\|_{L_{\infty}} + \|g_0(\widetilde{\boldsymbol{f}}, x_j)\|_{L_{\infty}})^2 \mathbb{P}(\mathbf{H}^{\dagger} \widetilde{\boldsymbol{F}} \notin [-2b, 2b]^r) \\ & \lesssim (2M_{\infty})^2 \mathbb{P}(\|d^{-1}\mathbf{H}^{\dagger} \mathbf{W}^{\top} \boldsymbol{u}\|_2 \ge b) \le \mathbb{E}[\|d^{-1}\mathbf{H}^{\dagger} \mathbf{W}^{\top} \boldsymbol{u}\|_2^2] \lesssim \delta_{\mathbf{f}}^2. \end{split}$$

Combining the two equation displays above yields the conclusion of the lemma.

Lemma B.4 (Bias for approximating the Riesz representer). Under Assumptions 3.1 to 3.4 and B.1 to B.4,

$$\left[\int \{ \widetilde{\alpha}_n(\widetilde{\boldsymbol{f}}, x_j) - \alpha_n^{\text{null}}(\widetilde{\boldsymbol{f}}, x_j) \}^2 d\mathbb{P} \right]^{1/2} \lesssim \nu_{\mathbf{a}, n} + \delta_{\mathbf{f}}.$$

Proof. As usual, let \widetilde{F}_k denote the k-th element of \widetilde{F} and moreover let $[\mathbf{W}^{\top}]_{k}$ denote the k-th row of \mathbf{W}^{\top} . Using the decomposition (B.4),

$$|\widetilde{F}_{k}| \leq |\mathbf{H}_{k} \cdot \boldsymbol{F}| + |d^{-1}[\mathbf{W}^{\top}]_{k} \cdot \boldsymbol{u}| \leq |\mathbf{H}_{k} \cdot \boldsymbol{F}| + ||d^{-1}\mathbf{W}^{\top}\boldsymbol{u}||_{2}.$$
 (B.9)

By Assumption 3.4, the first term $|\mathbf{H}_{k}.\mathbf{F}|$ in the last step above is uniformly bounded over $k \in \{1, \ldots, \overline{r}\}$ by $c_b b/2$ with probability $1 - \delta_{\mathrm{f}}^2$. The second term has been bounded in (B.8) and yields, by Markov's inequality,

$$\mathbb{P}(\|d^{-1}\mathbf{W}^{\top}\boldsymbol{u}\|_2 \ge c_b b/2) \lesssim \delta_{\mathrm{f}}^2$$

Hence, we conclude that with probability at least $1 - C\delta_{\rm f}^2$, $\tilde{F} \in [-c_b b, c_b b]^{\bar{r}}$. Then, using argument similar to that in the proof of Lemma B.3, the lemma holds because

$$\begin{split} & \int \{\widetilde{\alpha}_n(\widetilde{f}, x_j) - \alpha_n^{\text{null}}(\widetilde{f}, x_j)\}^2 d\mathbb{P} = \int \{\widetilde{\alpha}_n(\widetilde{f}, x_j) - \alpha_n^{\text{null}}(\widetilde{f}, x_j)\}^2 \mathbb{1}\{\widetilde{f} \in [-c_b b, c_b b]^{\overline{r}}\} d\mathbb{P} \\ & + \int \{\widetilde{\alpha}_n(\widetilde{f}, x_j) - \alpha_n^{\text{null}}(\widetilde{f}, x_j)\}^2 \mathbb{1}\{\widetilde{f} \notin [-c_b b, c_b b]^{\overline{r}}\} d\mathbb{P} \lesssim \nu_{\text{a},n}^2 + \delta_{\text{f}}^2. \end{split}$$

B.3 The stability of conditional expectation in Assumption **B.2**

This subsection addresses the issue that while model (2) yields the proper centering of the noise, namely $\mathbb{E}[\epsilon|\mathbf{F}, X_j] = 0$, after replacing \mathbf{F} by $\tilde{\mathbf{F}}$ we may no longer have $\mathbb{E}[\epsilon|\tilde{\mathbf{F}}, X_j] = 0$; instead, what is properly centered here is the adjusted noise term $\tilde{\epsilon} \equiv Y - \mathbb{E}[Y|\tilde{\mathbf{F}}, X_j]$. Naturally we will need to handle the difference $|\epsilon - \tilde{\epsilon}| = |\{Y - m_0(\mathbf{F}, X_j)\} - \{Y - \mathbb{E}[Y|\tilde{\mathbf{F}}, X_j]\}| = |\mathbb{E}[Y|\mathbf{F}, X_j] - \mathbb{E}[Y|\tilde{\mathbf{F}}, X_j]|$, whose L_2 norm is precisely the left-hand side of (B.1). This difference concerns the closeness between the two conditional expectations $\mathbb{E}[Y|\tilde{\mathbf{F}}, X_j]$ and $m_0(\mathbf{F}, X_j) = \mathbb{E}[Y|\mathbf{F}, X_j]$ under a perturbation of the conditioning variables. A simple version of such a "stability of conditional expectation" result appears in Bryc and Smolenski (1992); however, only one-dimensional random variables were examined and no explicit convergence rate was provided. We will provide in Lemma B.5 such a stability result which is more appropriate for our context.

For brevity, define $m^*(\mathbf{F}, \widetilde{\mathbf{F}}, X_j) = \mathbb{E}[Y|\mathbf{F}, \widetilde{\mathbf{F}}, X_j]$. Compared to the full regression function $m_0^*(\mathbf{F}, \mathbf{X}) = \mathbb{E}[Y|\mathbf{F}, \mathbf{X}]$ defined below (4), $m^*(\mathbf{F}, \widetilde{\mathbf{F}}, X_j)$ is obtained by conditioning

on a finer/smoother σ -algebra (because the σ -algebra generated by $(\boldsymbol{F}, \boldsymbol{\tilde{F}}, X_j)$ is contained in the σ -algebra generated by $(\boldsymbol{F}, \boldsymbol{X})$ and in fact $m^*(\boldsymbol{F}, \boldsymbol{\tilde{F}}, X_j) = \mathbb{E}[m_0^*(\boldsymbol{F}, \boldsymbol{X})|\boldsymbol{F}, \boldsymbol{\tilde{F}}, X_j]$. In turn, both $m_0(\boldsymbol{F}, X_j)$ and $\mathbb{E}[Y|\boldsymbol{\tilde{F}}, X_j]$ involved in (B.1) can be obtained by further smoothing $m^*(\boldsymbol{F}, \boldsymbol{\tilde{F}}, X_j)$ (as the proof of Lemma B.5 will show). Define $\boldsymbol{\Xi} = \mathbf{W}^\top \boldsymbol{u} \in \mathbb{R}^{\overline{r}}$. Further define the random vector $\mathbf{u}_{\boldsymbol{\Xi}, \boldsymbol{z}, d} = C_{\mathrm{f}}^{-1} \delta_{\mathrm{f}}^{-1} (\mathbf{H}^{\dagger} \boldsymbol{\Xi} - \mathbf{H}^{\dagger} \boldsymbol{z})$ where C_{f} is the constant in condition (v) in Lemma B.5, and the random functions g_{n} and g_{d} as, for $a \in \mathbb{R}$,

$$g_{n}(a) = \int m^{*}(\boldsymbol{F} + a\boldsymbol{u}_{\Xi,\boldsymbol{z},d}, \boldsymbol{H}\boldsymbol{F} + \frac{1}{d}\boldsymbol{z}, X_{j}) p_{\boldsymbol{F}}(\boldsymbol{F} + a\boldsymbol{u}_{\Xi,\boldsymbol{z},d}) p_{\Xi}(\boldsymbol{z}) p_{\mu_{j}}(\mu_{j} - \boldsymbol{B}_{j.}a\boldsymbol{u}_{\Xi,\boldsymbol{z},d}) d\boldsymbol{z},$$

$$g_{d}(a) = \int p_{\boldsymbol{F}}(\boldsymbol{F} + a\boldsymbol{u}_{\Xi,\boldsymbol{z},d}) p_{\Xi}(\boldsymbol{z}) p_{\mu_{j}}(\mu_{j} - \boldsymbol{B}_{j.}a\boldsymbol{u}_{\Xi,\boldsymbol{z},d}) d\boldsymbol{z}.$$

The reason we introduce these functions is because, in terms of them, by the derivations within the proof of Lemma B.5,

$$\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \widetilde{\boldsymbol{F}}, X_j] = \frac{g_n(C_f \delta_f)}{g_d(C_f \delta_f)}, \quad \mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \boldsymbol{F}, X_j] = \frac{g_n(0)}{g_d(0)}$$

hence, the left-hand side of (B.1) can be expressed with the help of the mean value theorem applied to the expressions above.

Lemma B.5 (Conditions for the validity of inequality (B.1)). Suppose that Assumptions 3.1 to 3.3, and B.2 excluding (B.1), hold. Moreover, suppose that

- (i) The function $m^* = m^*(\mathbf{f}, \widetilde{\mathbf{f}}, x_j)$ where $\mathbf{f} \in \mathbb{R}^r$ and $\widetilde{\mathbf{f}} \in \mathbb{R}^{\overline{r}}$ is bounded, and m^* admits a bounded gradient in the first argument \mathbf{f} which we denote by $\nabla m^*(\mathbf{f}, \widetilde{\mathbf{f}}, x_j)$;
- (ii) \mathbf{F} admits a density $p_{\mathbf{F}}$ with gradient $\nabla p_{\mathbf{F}}$;
- (iii) \boldsymbol{u} and \boldsymbol{F} are independent;
- (iv) Ξ and μ_j admits a joint density $p_{\Xi,\mu_j} = p_{\Xi,\mu_j}(\boldsymbol{\xi}, u_j)$, and p_{Ξ,μ_j} admits a partial derivative with respect to u_j which we denote by \dot{p}_{Ξ,μ_j} ;
- (v) The following bounds hold:

$$\mathbb{E}\left[\sup_{0 < a < C_{\mathrm{f}}\delta_{\mathrm{f}}}\left\{\frac{\int \nabla p_{F}(F + a\mathbf{u}_{\Xi,z,d})^{\top} p_{\Xi,\mu_{j}}(z,\mu_{j} - \mathbf{B}_{j}.a\mathbf{u}_{\Xi,z,d})\mathbf{u}_{\Xi,z,d}\mathrm{d}z}{g_{\mathrm{d}}(a)}\right\}^{2}\right] < \infty,$$
$$\mathbb{E}\left[\sup_{0 < a < C_{\mathrm{f}}\delta_{\mathrm{f}}}\left\{\frac{\int p_{F}(F + a\mathbf{u}_{\Xi,z,d})\dot{p}_{\Xi,\mu_{j}}(z,\mu_{j} - \mathbf{B}_{j}.a\mathbf{u}_{\Xi,z,d})\mathbf{B}_{j}.\mathbf{u}_{\Xi,z,d}\mathrm{d}z}{g_{\mathrm{d}}(a)}\right\}^{2}\right] < \infty;$$

(vi) The random vector Ξ satisfies $\|\mathbf{H}^{\dagger}\Xi\|_{2} \leq C_{f}\delta_{f}$ almost surely for a constant C_{f} . Then Inequality (B.1) holds.

Some remarks about the conditions in Lemma B.5 are in order. First, condition (vi) is not directly used in the proof of Lemma B.5 per se, but it facilitates our subsequent discussion about condition (vi). In any case, the proof of Lemma B.2 shows that $\mathbb{E} \| \mathbf{H}^{\dagger} \mathbf{\Xi} \|_2 \lesssim \delta_{\mathrm{f}}$, so condition (vi) is reasonable. Condition (vi) may appear complex at first; however, this condition simplifies vastly if we simply set a = 0. Take the second inequality in

condition (vi) as an example. When a = 0, the expectation on the left-hand side reduces to

$$\mathbb{E}\left[\left\{\frac{\int \dot{p}_{\Xi,\mu_j}(\boldsymbol{z},\mu_j)\mathbf{B}_{j\cdot}\mathbf{u}_{\Xi,\boldsymbol{z},d}\mathrm{d}\boldsymbol{z}}{\int p_{\Xi,\mu_j}(\boldsymbol{z},\mu_j)\mathrm{d}\boldsymbol{z}}\right\}^2\right] \stackrel{\text{condition }(\boldsymbol{v}\boldsymbol{i})}{\lesssim} \mathbb{E}\left[\left\{\frac{\int \dot{p}_{\Xi,\mu_j}(\boldsymbol{z},\mu_j)\mathrm{d}\boldsymbol{z}}{\int p_{\Xi,\mu_j}(\boldsymbol{z},\mu_j)\mathrm{d}\boldsymbol{z}}\right\}^2\right] = \mathbb{E}\left[\left\{\frac{\dot{p}_{\mu_j}(\mu_j)}{p_{\mu_j}(\mu_j)}\right\}^2\right]$$

where \dot{p}_{μ_j} denotes the derivative of the marginal density p_{μ_j} of μ_j and in the last step we have exchanged integration and differentiation. We recognize the rightmost term above as simply the Fisher information for location of the idiosyncratic term μ_j , and the boundedness of this information is a very weak condition. Thus, condition (vi) can be seen as a slightly strengthened version of the boundededness of this Fisher information under a small perturbation by the random variable $a\mathbf{u}_{\Xi,z,d}$ where $0 < a < C_f \delta_f$.

Proof of Lemma B.5. By the smoothing property of conditional expectations (see, for instance, property (10) on p. 348 of Resnick (2005)), $\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \boldsymbol{F}, X_j] = \mathbb{E}[\mathbb{E}[Y|\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j] | \boldsymbol{F}, X_j] = \mathbb{E}[Y|\boldsymbol{F}, X_j] = m_0(\boldsymbol{F}, X_j)$, and similarly $\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \widetilde{\boldsymbol{F}}, X_j] = \mathbb{E}[Y|\widetilde{\boldsymbol{F}}, X_j]$. Thus,

$$\mathbb{E}[\{\mathbb{E}[Y|\widetilde{\boldsymbol{F}},X_j] - m_0(\boldsymbol{F},X_j)\}^2] = \mathbb{E}[\{\mathbb{E}[m^*(\boldsymbol{F},\widetilde{\boldsymbol{F}},X_j)|\widetilde{\boldsymbol{F}},X_j] - \mathbb{E}[m^*(\boldsymbol{F},\widetilde{\boldsymbol{F}},X_j)|\boldsymbol{F},X_j]\}^2].$$

We now give explicit expressions for $\mathbb{E}[m^*(\boldsymbol{F}, \boldsymbol{\widetilde{F}}, X_j) | \boldsymbol{\widetilde{F}}, X_j]$ and $\mathbb{E}[m^*(\boldsymbol{F}, \boldsymbol{\widetilde{F}}, X_j) | \boldsymbol{F}, X_j]$. We start from first of these, namely

$$\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \widetilde{\boldsymbol{F}}, X_j] = \frac{\int m^*(\boldsymbol{f}, \widetilde{\boldsymbol{F}}, X_j) p_{\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j}(\boldsymbol{f}, \widetilde{\boldsymbol{F}}, X_j) \mathrm{d}\boldsymbol{f}}{p_{\widetilde{\boldsymbol{F}}, X_j}(\widetilde{\boldsymbol{F}}, X_j)}, \quad (B.10)$$

where $p_{\tilde{F},X_j}$ and p_{F,\tilde{F},X_j} denote respectively the joint densities of \tilde{F} and X_j , and of F, \tilde{F} and X_j . Recalling the decomposition (B.4) which gives $\tilde{F} = \mathbf{H}F + \frac{1}{d}\Xi$, and letting \mathbf{B}_j . denote the *j*-th row of \mathbf{B} and so $X_j = \mathbf{B}_j \cdot F + \mu_j$, we note that for $\tilde{a} \in \mathbb{R}^{\bar{r}}$ and $b \in \mathbb{R}$, that $(\tilde{F}, X_j) = (\mathbf{H}F + \frac{1}{d}\Xi, \mathbf{B}_j \cdot F + \mu_j) = (\tilde{a}^{\top}, b)^{\top}$ is equivalent to $\Xi = d(\tilde{a} - \mathbf{H}F)$ and $\mu_j = b - \mathbf{B}_j \cdot F = b - \mathbf{B}_j \cdot \mathbf{H}^{\dagger} (\tilde{F} - \frac{1}{d}\Xi)$. Then, by the usual convolution argument, the joint density of $p_{\tilde{F},X_j}$ is given by

$$\begin{split} p_{\widetilde{F},X_{j}}(\widetilde{\boldsymbol{a}},b) &= \frac{\mathrm{d}^{2}}{\mathrm{d}\widetilde{\boldsymbol{a}}\mathrm{d}\boldsymbol{b}} \mathbb{P}(\widetilde{F} \leqslant \widetilde{\boldsymbol{a}},X_{j} \leqslant b) = \frac{\mathrm{d}^{2}}{\mathrm{d}\widetilde{\boldsymbol{a}}\mathrm{d}\boldsymbol{b}} \mathbb{P}(\mathbf{H}\boldsymbol{F} + \frac{1}{d}\boldsymbol{\Xi} \leqslant \widetilde{\boldsymbol{a}},\mathbf{B}_{j}.\boldsymbol{F} + \mu_{j} \leqslant b) \\ &= \frac{\mathrm{d}^{2}}{\mathrm{d}\widetilde{\boldsymbol{a}}\mathrm{d}\boldsymbol{b}} \mathbb{P}(\boldsymbol{\Xi} \leqslant d(\widetilde{\boldsymbol{a}} - \mathbf{H}\boldsymbol{F}),\mu_{j} \leqslant b - \mathbf{B}_{j}.\boldsymbol{F}) = \frac{\mathrm{d}^{2}}{\mathrm{d}\widetilde{\boldsymbol{a}}\mathrm{d}\boldsymbol{b}} \int^{d(\widetilde{\boldsymbol{a}} - \mathbf{H}\boldsymbol{f})} \int^{b - \mathbf{B}_{j}.\boldsymbol{f}} p_{\boldsymbol{F}}(\boldsymbol{f}) p_{\boldsymbol{\Xi},\mu_{j}}(\boldsymbol{w},v) \mathrm{d}v \mathrm{d}\boldsymbol{w} \mathrm{d}\boldsymbol{f} \\ &= d^{\overline{r}} \int p_{\boldsymbol{F}}(\boldsymbol{f}) p_{\boldsymbol{\Xi},\mu_{j}}(d(\widetilde{\boldsymbol{a}} - \mathbf{H}\boldsymbol{f}), b - \mathbf{B}_{j}.\boldsymbol{f}) \mathrm{d}\boldsymbol{f} \end{split}$$

where the last step follows by the Leibniz integral rule. Now, $p_{\tilde{F},X_j}(\tilde{F},X_j)$ can serve as the denominator of $\mathbb{E}[m^*(F,\tilde{F},X_j)|\tilde{F},X_j]$ in (B.10). Analogously, in the numerator of (B.10), for $\boldsymbol{a} \in \mathbb{R}^r$,

$$p_{\boldsymbol{F},\tilde{\boldsymbol{F}},X_j}(\boldsymbol{a},\tilde{\boldsymbol{a}},b) = d^{\overline{r}} p_{\boldsymbol{F}}(\boldsymbol{a}) p_{\boldsymbol{\Xi},\mu_j}(d(\tilde{\boldsymbol{a}}-\mathbf{H}\boldsymbol{a}),b-\mathbf{B}_j,\boldsymbol{a}).$$

Consequently,

 $\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \widetilde{\boldsymbol{F}}, X_j]$

$$= \frac{\int m^*(\boldsymbol{f}, \widetilde{\boldsymbol{F}}, X_j) p_{\boldsymbol{F}}(\boldsymbol{f}) p_{\Xi,\mu_j}(d(\widetilde{\boldsymbol{F}} - \mathbf{H}\boldsymbol{f}), X_j - \mathbf{B}_j \cdot \boldsymbol{f}) d\boldsymbol{f}}{\int p_{\boldsymbol{F}}(\boldsymbol{f}) p_{\Xi,\mu_j}(d(\widetilde{\boldsymbol{F}} - \mathbf{H}\boldsymbol{f}), X_j - \mathbf{B}_j \cdot \boldsymbol{f}) d\boldsymbol{f}}$$

$$= \frac{\left(\int m^*(\boldsymbol{F} + \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{\Xi} - \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{z}, \mathbf{H}\boldsymbol{F} + \frac{1}{d}\boldsymbol{z}, X_j) p_{\boldsymbol{F}}(\boldsymbol{F} + \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{\Xi} - \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{z})}{\times p_{\Xi,\mu_j}(\boldsymbol{z},\mu_j - \mathbf{B}_j \cdot (\frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{\Xi} - \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{z})) d\boldsymbol{z}}\right)}{\int p_{\boldsymbol{F}}(\boldsymbol{F} + \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{\Xi} - \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{z}) p_{\Xi,\mu_j}(\boldsymbol{z},\mu_j - \mathbf{B}_j \cdot (\frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{\Xi} - \frac{1}{d}\mathbf{H}^{\dagger}\boldsymbol{z})) d\boldsymbol{z}} = \frac{g_n(C_f\delta_f)}{g_d(C_f\delta_f)}$$

where in the transition to the second to last step we have conducted a change of the integration variable from \boldsymbol{f} to $\boldsymbol{z} = d(\tilde{\boldsymbol{F}} - \mathbf{H}\boldsymbol{f})$. Meanwhile, for the second conditional expectation $\mathbb{E}[m^*(\boldsymbol{F}, \tilde{\boldsymbol{F}}, X_j) | \boldsymbol{F}, X_j]$, we can analogously obtain

$$\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j) | \boldsymbol{F}, X_j] = \frac{\int m^*(\boldsymbol{F}, \mathbf{H}\boldsymbol{F} + \frac{1}{d}\boldsymbol{z}, X_j) p_{\boldsymbol{\Xi}, \mu_j}(\boldsymbol{z}, \mu_j) d\boldsymbol{z}}{p_{\mu_j}(\mu_j)} = \frac{g_n(0)}{g_d(0)}$$

Next, denote by \dot{g}_n and \dot{g}_d the derivatives of g_n and g_d respectively. By the mean value theorem, for some random \tilde{a} strictly between 0 and $C_f \delta_f$,

$$\begin{aligned} &|\mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j)|\widetilde{\boldsymbol{F}}, X_j] - \mathbb{E}[m^*(\boldsymbol{F}, \widetilde{\boldsymbol{F}}, X_j)|\boldsymbol{F}, X_j]| = \left|\frac{g_n(C_{\mathrm{f}}\delta_{\mathrm{f}})}{g_{\mathrm{d}}(C_{\mathrm{f}}\delta_{\mathrm{f}})} - \frac{g_n(0)}{g_{\mathrm{d}}(0)}\right| \\ &= \left|\left\{\frac{\dot{g}_n(\widetilde{a})}{g_{\mathrm{d}}(\widetilde{a})} - \frac{g_n(\widetilde{a})}{g_{\mathrm{d}}(\widetilde{a})}\frac{\dot{g}_{\mathrm{d}}(\widetilde{a})}{g_{\mathrm{d}}(\widetilde{a})}\right\}\right| C_{\mathrm{f}}\delta_{\mathrm{f}} \overset{\text{condition }(i)}{\lesssim} \left\{\left|\frac{\dot{g}_n(\widetilde{a})}{g_{\mathrm{d}}(\widetilde{a})}\right| + \left|\frac{\dot{g}_{\mathrm{d}}(\widetilde{a})}{g_{\mathrm{d}}(\widetilde{a})}\right|\right\} C_{\mathrm{f}}\delta_{\mathrm{f}}.\end{aligned}$$

Finally, Inequality (B.1) follows by applying condition (vi) to the last term above.

B.4 Useful bounds on key stochastic terms

We first present a couple of results that will be used repeatedly for bounding some key stochastic terms throughout our proofs, for instance in the proof of Theorem 3.1 in Section C.1. Here, the key element is to utilize sharp weighted empirical process techniques in, for instance, Giné et al. (2003) to obtain, for each indexing function (for instance θ on the lefthand sides of (B.14) and (B.15)), a bound on the empirical sum of the index function that is related to the norm of that particular indexing function (for instance, $\|\theta(Y_i, F_i, X_i)\|_{L_2}$ on the right-hand sides of (B.14) and (B.15)). Bounds of this kind become sharper when the norm of the indexing function itself is small, and are key to obtaining good rates for \hat{g}_n and our subsequent derivative estimators.

Define the quantity $\nu_{s,n}$ that will feature prominently in the our subsequent Lemma B.6 and Proposition B.7:

$$\nu_{\mathrm{s},n} = \left\{ \frac{V_{\mathcal{F}_n(\overline{r}+1)}\log(n)}{n} \right\}^{1/2}.$$
(B.11)

Lemma B.6 (Lemma 2 in Fan and Gu (2023)). Let \mathcal{G} be a class of uniformly bounded functions on \mathbb{R}^{1+r+d} , for which we can choose a constant (bounded) envelope function G. Assume that the covering number for \mathcal{G} admits the bound (A.1) with $v = V_{\mathcal{F}_n(\bar{r}+1)}$, that is, the class \mathcal{G} is of VC-type (Def. A.4) with index $V_{\mathcal{F}_n(\bar{r}+1)}$. Then, there exist some universal constants c_1 , c_2 , c_3 such that for all $\sqrt{s'} \ge c_1 \sqrt{n\nu_{s,n}} / {\zeta(1-\zeta)}$ with $0 < \zeta < 1$, on an event $\mathcal{A}_{1,\zeta,s'}$ satisfying

$$\mathbb{P}(\mathcal{A}_{1,\zeta,s'}) \ge 1 - c_2 \exp(-c_3 \zeta^2 (1-\zeta)s') \tag{B.12}$$

we have

$$\forall g \in \mathcal{G}, \quad \left|\frac{1}{n}\sum_{i=1}^{n} \{g^2(Y_i, \boldsymbol{F}_i, \boldsymbol{X}_i) - \mathbb{E}[g^2(Y, \boldsymbol{F}, \boldsymbol{X})]\}\right| \leq \zeta \left\{ \int g^2(Y_i, \boldsymbol{F}_i, \boldsymbol{X}_i) d\mathbb{P} + \frac{s'}{n} \right\}.$$
(B.13)

Proof. As hinted in the lemma statement, this is essentially Lemma 2 in Fan and Gu (2023) with minor notational differences. Alternatively, one could rely on Lemma 2 in Giné et al. (2003). We omit the proof. \Box

Recall the adjusted noise $\tilde{\epsilon} = Y - \mathbb{E}[Y|\tilde{F}, X_j]$ from Section B.3. Then, define $\tilde{\epsilon}_i = Y_i - \mathbb{E}[Y_i|\tilde{F}_i, X_{i,j}], i = 1, ..., n$, which form i.i.d. copies of $\tilde{\epsilon}$.

Proposition B.7. Let \mathcal{G} be a class of uniformly bounded functions on $\mathbb{R}^{\overline{r}+1}$, for which we can choose a constant (bounded) envelope function G. Assume that the covering number for \mathcal{G} admits the bound (A.1) with $v = V_{\mathcal{F}_n(\overline{r}+1)}$. Then, on an event $\mathcal{A}_{2,s}$ with $\mathbb{P}(\mathcal{A}_{2,s}) \geq 1 - 12e^{-s}$, for a constant c_4 ,

$$\forall \theta \in \mathcal{G}, \quad \left|\frac{1}{n} \sum_{i=1}^{n} \widetilde{\epsilon}_{i} \theta(\widetilde{F}_{i}, X_{i,j})\right| \leq c_{4} \left\{ \|\theta(\widetilde{F}, X_{j})\|_{L_{2}} + \nu_{\mathrm{s},n} \right\} \left\{ \nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}} \right\}.$$
(B.14)

Moreover, instead let \mathcal{G} be a class of uniformly bounded functions on \mathbb{R}^{1+r+d} , again with a covering number satisfying the bound (A.1) with $v = V_{\mathcal{F}_n(\bar{r}+1)}$. Then on an event $\mathcal{A}_{3,s}$ with $\mathbb{P}(\mathcal{A}_{3,s}) \ge 1 - 12e^{-s}$, for a constant c_5 ,

$$\forall \theta \in \mathcal{G}, \quad \left| \frac{1}{n} \sum_{i=1}^{n} \left\{ \theta(Y_i, \boldsymbol{F}_i, \boldsymbol{X}_i) - \mathbb{E}[\theta(Y_i, \boldsymbol{F}_i, \boldsymbol{X}_i)] \right\} \right| \leq c_5 \left\{ \|\theta(Y, \boldsymbol{F}, \boldsymbol{X})\|_{L_2} + \nu_{\mathrm{s}, n} \right\} \left\{ \nu_{\mathrm{s}, n} + \sqrt{\frac{s}{n}} \right\}$$
(B.15)

Proof. We first prove (B.14). Our main strategy is to apply the proof of Theorem 4 in Giné et al. (2003). For simplicity we assume that the uniform bound on the magnitudes of functions in \mathcal{G} is one, so we could simply choose the constant envelope G = 1.

We let $s = s_n = 1$, $l = l_n = \lceil \log(\nu_{s,n}^{-1}) \rceil$ (where " $\lceil \cdot \rceil$ " is the ceiling operator that returns the smallest integer no smaller than the argument), $r = r_n = e^{-l_n} \in [e^{-1}\nu_{s,n}, \nu_{s,n}]$, and $\rho_j = r_n e^j$. Define the function classes $\mathcal{G}(a) = \{\theta : \theta \in \mathcal{G}, \|\theta\|_{L_2} \leq a\}$, $\mathcal{G}(a, b] = \mathcal{G}(b) \setminus \mathcal{G}(a)$, and $\mathcal{G}_e(a) = \mathcal{G}(a/e, a]$.

Our general proof strategy based on the aforementioned Theorem 4 in Giné et al. (2003) is to control the stochastic behavior of the empirical process $\sum_{i=1}^{n} \tilde{\epsilon}_{i} \theta(\tilde{F}_{i}, X_{i,j})$ within each "shell" $\theta \in \mathcal{G}_{e}(\rho_{j})$ for the shells numbered $j = 1, \ldots, l$.

As our first step, we control the tail probability with respect to the expectation of the (supremum of the) empirical process within each shell. Recall the sub-Gaussian norm $C_{\epsilon} < 0$ of ϵ from Assumption B.1. Also note that the bound M_{∞} on the full conditional expectation $m_0^*(\mathbf{F}, \mathbf{X})$ implies the same bound on any of its smoothed versions including $m_0(\mathbf{F}, X_j) = \mathbb{E}[m_0^*(\mathbf{F}, \mathbf{X})|\mathbf{F}, X_j]$ and $\mathbb{E}[Y|\widetilde{\mathbf{F}}, X_j] = \mathbb{E}[m_0^*(\mathbf{F}, \mathbf{X})|\widetilde{\mathbf{F}}, X_j]$ due to the monotonicity property of conditional expectations. Thus, $\tilde{\epsilon} - \epsilon = m_0(\mathbf{F}, X_j) - \mathbb{E}[Y|\widetilde{\mathbf{F}}, X_j]$ is bounded, and so $\tilde{\epsilon}$ is sub-Gaussian as well and we assume its sub-Gaussian norm is a constant \tilde{C}_{ϵ} (with $\tilde{C}_{\epsilon} < \infty$). By first Theorem 4 and then Eq. (13) in Adamczak (2008), we know that

$$\mathbb{P}\left(\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\left|\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{\boldsymbol{F}}_{i},X_{i,j})\right| \geq 2\mathbb{E}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\left|\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{\boldsymbol{F}}_{i},X_{i,j})\right| + s\right) \\
\leq \exp\left(-\frac{s^{2}}{4n\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\|\widetilde{\epsilon}\theta(\widetilde{\boldsymbol{F}},X_{j})\|_{L_{2}}^{2}}\right) + 3\exp\left(-\frac{s^{2}}{C^{2}\|\max_{i\in\{1,\dots,n\}}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}|\widetilde{\epsilon}_{i}\theta(\widetilde{\boldsymbol{F}}_{i},X_{i,j})|\|_{\psi_{2}}^{2}}\right)$$

$$\leq \exp\left(-\frac{s^2}{4n\|\widetilde{\epsilon}\|_{L_2}^2\rho_j^2}\right) + 3\exp\left(-\frac{s^2}{C^2\max_{i\in\{1,\dots,n\}}\|\sup_{\theta\in\mathcal{G}_e(\rho_j)}\widetilde{\epsilon}_i\theta(\widetilde{F}_i,X_{i,j})\|_{\psi_2}^2\log(n)}\right)$$
$$\leq \exp\left(-\frac{s^2}{4n\|\widetilde{\epsilon}\|_{L_2}^2\rho_j^2}\right) + 3\exp\left(-\frac{s^2}{C^2\widetilde{C}_{\epsilon}^2\log(n)}\right).$$

Then, absorbing \widetilde{C}_{ϵ} into C and setting $s = \{2 \| \widetilde{\epsilon} \|_{L_2} \rho_j \sqrt{n} + C \log^{1/2}(n) \} \sqrt{u}$ yields

$$\mathbb{P}\left(\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\left|\frac{1}{n}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{F}_{i},X_{i,j})\right| \leq 2\mathbb{E}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\left|\frac{1}{n}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{F}_{i},X_{i,j})\right| + \frac{2}{\sqrt{n}}\|\widetilde{\epsilon}\|_{L_{2}}\rho_{j}\sqrt{u} + C\frac{\log^{1/2}(n)}{n}\sqrt{u}\right) \ge 1 - 4e^{-u}. \quad (B.16)$$

Then, conforming to the notation in the proof of Theorem 4 in Giné et al. (2003), define the events

$$\mathcal{E}_{j}^{+}(s) = \left\{ \sup_{\theta \in \mathcal{G}_{e}(\rho_{j})} \left| \frac{1}{n} \sum_{i=1}^{n} \widetilde{\epsilon}_{i} \theta(\widetilde{F}_{i}, X_{i,j}) \right| \leq 2\mathbb{E} \sup_{\theta \in \mathcal{G}_{e}(\rho_{j})} \left| \frac{1}{n} \sum_{i=1}^{n} \widetilde{\epsilon}_{i} \theta(\widetilde{F}_{i}, X_{i,j}) \right| \\ + \frac{2\sqrt{s+2\log(l-j+1)}}{\sqrt{n}} \|\widetilde{\epsilon}\|_{L_{2}} \rho_{j} + C \frac{\log^{1/2}(n)}{n} \sqrt{s+2\log(l-j+1)} \right\}.$$

Then $\mathbb{P}(\cap_{j=1}^{l} \mathcal{E}_{j}^{+}(s)) \ge 1 - 8e^{-s}.$

As our second step, with the tail probability on hand, we shall bound the expectation $\mathbb{E}\sup_{\theta\in\mathcal{G}_e(\rho_j)} \left|\sum_{i=1}^n \tilde{\epsilon}_i\theta(\tilde{F}_i, X_{i,j})\right|$ within each shell. Further introduce the notation that $\omega_n(a) = n^{-1/2}\mathbb{E}\sup_{\theta\in\mathcal{G}_e(\rho_j)} \left|\sum_{i=1}^n \tilde{\epsilon}_i\theta(\tilde{F}_i, X_{i,j})\right|$ if $a \in (\rho_{j-1}, \rho_j]$. We now find a function $\omega : [0,1] \to \mathbb{R}^+$ that satisfies the conditions in Theorem 4 in Giné et al. (2003), that is, (i) ω is non-decreasing; (ii) $\omega(a)/a$ is non-increasing; (iii) $\omega_n(a) \leq \omega(a)$, $\forall a \in [r_n, s_n]$; (iv) $\sup_{a \in (0,1]} a \sqrt{\log \log(1/a)}/\omega(a) \leq K < \infty$. We set out to find such a ω .

Note that we can regard $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} (\tilde{\epsilon}_i/\tilde{C}_{\epsilon}) \theta(\tilde{F}_i, X_{i,j})$, (indexed by) $\theta \in \mathcal{G}_e(\rho_j)$ as the empirical process $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \theta(\tilde{F}_i, X_{i,j}), \theta \in \mathcal{G}_e(\rho_j)$ symmetrized by the random variables $\tilde{\epsilon}_i/\tilde{C}_{\epsilon}$, $i \in \{1, \ldots, n\}$. Although the $\tilde{\epsilon}_i/\tilde{C}_{\epsilon}$'s are not Rademacher random variables, by our earlier argument, when conditioning on the \tilde{F}_i 's and $X_{i,j}$'s, the process $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} (\tilde{\epsilon}_i/\tilde{C}_{\epsilon}) \theta(\tilde{F}_i, X_{i,j})$ is nevertheless centered and sub-Gaussian with respect to the (random) distance d on $\mathcal{G}_e(\rho_j)$ defined as $d^2(\theta_1, \theta_2) = \frac{1}{n} \sum_{i=1}^{n} (\theta_1 - \theta_2)^2 (\tilde{F}_i, X_{i,j}) = \|\theta_1 - \theta_2\|_{L_2(\mathbb{P}_n)}^2$. Thus, the proof of Theorem 3.5.4 in Giné and Nickl (2016), in particular the part that invokes Theorem 3.5.1, applies. Specifically, let Q be a generic measure on (\tilde{F}, X_j) , and recall that G = 1 is the envelope for the class $\mathcal{G}_e(\rho_j)$ (so simply $\|G\|_{L_2(Q)} = 1$, though at times we will retain the dependence on $\|G\|_{L_2(Q)}$ to conform to the notations in van der Vaart and Wellner (1996) and Giné and Nickl (2016)), let $N(\mathcal{G}_e(\rho_j), L_2(Q), \tau)$ be the covering number for the class $\mathcal{G}_e(\rho_j)$ by balls of radius τ in the $L_2(Q)$ norm, and let the constants A_1, A_2 be as in Theorem 3.5.4 in Giné and Nickl (2016). We then obtain from that theorem that, if $a \in (\rho_{j-1}, \rho_j]$,

$$\frac{1}{\widetilde{C}_{\epsilon}}\omega_{n}(a) = \mathbb{E}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\left|\frac{1}{\sqrt{n}}\sum_{i=1}^{n}(\widetilde{\epsilon}_{i}/\widetilde{C}_{\epsilon})\theta(\widetilde{F}_{i},X_{i,j})\right| \\ \leqslant A_{1}\int_{0}^{\rho_{j}}\sup_{Q}\sqrt{\log\{2N(\mathcal{G}_{e}(\rho_{j}),L_{2}(Q),\tau\|G\|_{L_{2}(Q)})\}}d\tau$$

+
$$A_2 \left\{ \int_0^{\rho_j} \sup_Q \sqrt{\log\{2N(\mathcal{G}_e(\rho_j), L_2(Q), \tau \|G\|_{L_2(Q)})\}} d\tau \right\}^2 / \sqrt{n\rho_j^2}$$
.

By Theorem 2.6.7 in van der Vaart and Wellner (1996), we can take, for a large enough constant C_1 , for \mathcal{Q} the set of all generic measures Q on (\widetilde{F}, X_j) and for all $0 < \tau \leq 1$,

$$\sup_{Q \in \mathcal{Q}} \log\{2N(\mathcal{G}_{e}(\rho_{j}), L_{2}(Q), \tau \| G \|_{L_{2}(Q)})\} \leq C_{1} V_{\mathcal{F}_{n}(\bar{r}+1)} \log(2/\tau).$$
(B.17)

and thus, for a large enough constant C_2 ,

$$\omega(u) = C_2 \left\{ V_{\mathcal{F}_n(\bar{r}+1)}^{1/2} \log^{1/2} (2/u) u + n^{-1/2} V_{\mathcal{F}_n(\bar{r}+1)} \log(2/u) \right\}.$$

Finally, we put everything together. Now, on the event $\cap_{j=1}^{l} \mathcal{E}_{j}^{+}(s), \forall 1 \leq j \leq l, \forall \theta \in \mathcal{G}_{e}(\rho_{j}),$

$$\left|\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{F}_{i},X_{i,j})\right| \leq 2\omega_{n}(\|\theta\|_{L_{2}}) + 2\sqrt{s+2\log(l-j+1)}\|\widetilde{\epsilon}\|_{L_{2}}\rho_{j} + C\frac{\log^{1/2}(n)}{\sqrt{n}}\sqrt{s+2\log(l-j+1)} \leq 2\omega(\rho_{j}) + (2\|\widetilde{\epsilon}\|_{L_{2}} + C)\sqrt{s+2\log(l-j+1)}\rho_{j}.$$

By our condition (iv) on ω and the derivations in the proof of Theorem 4 in Giné et al. (2003), for some positive constant C_3 ,

$$\max_{1 \le j \le l} \frac{\rho_j \sqrt{\log(l-j+1)}}{\omega(\rho_j)} \le C_3$$

Hence, on the same event (that is, the event $\cap_{j=1}^{l} \mathcal{E}_{j}^{+}(s)$), $\forall j, \forall \theta \in \mathcal{G}_{e}(\rho_{j})$,

$$\left| \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \widetilde{\epsilon}_{i} \theta(\widetilde{F}_{i}, X_{i,j}) \right| \leq \left\{ 2 + (2\|\widetilde{\epsilon}\|_{L_{2}} + C)\sqrt{2}C_{3} \right\} \omega(\rho_{j}) + (2\|\widetilde{\epsilon}\|_{L_{2}} + C)\rho_{j}\sqrt{s} \\ \leq C \left[\left\{ V_{\mathcal{F}_{n}(\overline{r}+1)}^{1/2} \log^{1/2}(n) + \sqrt{s} \right\} \|\theta\|_{L_{2}} + n^{-1/2}V_{\mathcal{F}_{n}(\overline{r}+1)} \log(n) \right].$$

for the class $\mathcal{G}(r_n)$, we can just invoke (B.16) but with $\mathcal{G}_e(\rho_j)$ replaced by $\mathcal{G}(r_n)$, to conclude that

$$\mathbb{P}\left[\sup_{\theta\in\mathcal{G}(r_n)}\left|\frac{1}{n}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\theta(\widetilde{F}_{i},X_{i,j})\right|\leqslant C\left\{\nu_{\mathrm{s},n}+\sqrt{\frac{s}{n}}\right\}\nu_{\mathrm{s},n}\right]\geqslant 1-4e^{-s}.$$

Combining the two cases above, we obtain the part of the conclusion of the proposition regarding (B.14). The other part of the proposition regarding (B.15) follows by minor modification. (Here the process on the left-hand side of (B.15) is properly centered for each θ , so we can proceed with the usual symmetrization by Rademacher random variables.) \Box

B.5 The stochastic term in the proof of Theorem 3.2

Recall the definition of α_n^{null} , which is the *t*-independent value of $\alpha_{t,n}^*$ under the null hypothesis, below (26) and $\tilde{\alpha}_n$ as in (B.3). We recall that the $(\cdot)_j^s$ as the smoothed *j*-th partial derivative of the argument function as in (10). Define, for $\alpha \in \mathcal{F}_n(\bar{r}+1)$, the function $\theta(\cdot; \alpha) : \mathbb{R}^{\bar{r}+1} \to \mathbb{R}$ indexed by α as

$$\theta(\widetilde{\boldsymbol{f}}, x_j; \alpha) = (\alpha - \widetilde{\alpha}_n)_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j) \mathbb{1}\{x_j \in \mathcal{B}_h\} - \{(\alpha - \widetilde{\alpha}_n)\alpha_n^{\mathrm{null}}\}(\widetilde{\boldsymbol{f}}, x_j).$$
(B.18)

Note that, under the null hypothesis or not, for each α and t, the function $\theta(\cdot; \alpha)$ is always centered under \mathbb{P} . Thus, it is not necessary to center $\mathbb{P}_n \theta(\cdot; \alpha)$ at $\mathbb{P}\theta(\cdot; \alpha) = 0$ in (B.20).

We show the following result that will be useful for bounding the major stochastic term in the proof of Theorem 3.2.

Proposition B.8. Assume that Assumptions 3.1 to 3.4 hold, and in addition the bandwidth h satisfies

$$h^2 V_{\mathcal{F}_n(\bar{r}+1)} \log(n) \ge 1. \tag{B.19}$$

Then on an event with probability at least $1 - 12e^{-s}$, for a constant c_6 ,

$$\forall \alpha \in \mathcal{F}_n(\overline{r}+1), \quad |\mathbb{P}_n\theta(\cdot;\alpha)| \leq c_6 \left[\{ \|\theta(\cdot;\alpha)\|_{L_2} + \nu_{\mathrm{s},n} \} \left\{ \nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}} \right\} + \frac{1}{h}\nu_{\mathrm{s},n}^2 \right]. \tag{B.20}$$

Proof. We mimic the proof of our Proposition B.7. (Alternatively, we could adapt the proof of Theorem 4 in Giné et al. (2003), which will improve some multiplicative constant.) Define the function class $\mathcal{G}_n = \{\theta(\cdot; \alpha) : \alpha \in \mathcal{F}_n(\overline{r}+1)\}$. By the form of the function θ , we can take $G = c_{\mathcal{G}}/h$ for a constant $c_{\mathcal{G}}$ large enough as the envelope for the class \mathcal{G}_n .

Similar to the proof of Proposition B.7 but with some minor modifications, we let $s = s_n = \|G\|_{L_2} = c_{\mathcal{G}}/h, \ l = l_n = [\log((c_{\mathcal{G}}/h)\nu_{s,n}^{-1})], \ r = r_n = s_n e^{-l_n} \in [e^{-1}\nu_{s,n}, \nu_{s,n}], \text{ and } \rho_j = r_n e^j$. Define the function classes $\mathcal{G}(a) = \{\theta \in \mathcal{G}_n, \|\theta\|_{L_2} \leq a\}, \ \mathcal{G}(a, b] = \mathcal{G}(b) \setminus \mathcal{G}(a), \text{ and } \mathcal{G}_e(a) = \mathcal{G}(a/e, a].$

By Theorem 4 in Adamczak (2008) with $\eta = 1$ (which is in fact allowed), $\delta = 1$ and $\alpha = 2$, we know that

$$\mathbb{P}\left(\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}|n\mathbb{P}_{n}\theta| \geq 2\mathbb{E}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}|n\mathbb{P}_{n}\theta| + s\right) \\
\leq \exp\left(-\frac{s^{2}}{4n\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}\|\theta(\widetilde{F},X_{j}) - \mathbb{E}\theta(\widetilde{F},X_{j})\|_{L_{2}}^{2}}\right) \\
+ 3\exp\left(-\frac{s^{2}}{C^{2}\|\max_{i\in\{1,\dots,n\}}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})}|\theta(\widetilde{F}_{i},X_{i,j}) - \mathbb{E}\theta(\widetilde{F},X_{i,j})|\|_{\psi_{2}}^{2}}\right) \\
\leq \exp\left(-\frac{s^{2}}{4n\rho_{j}^{2}}\right) + 3\exp\left(-\frac{s^{2}}{C^{2}c_{\mathcal{G}}^{2}/h^{2}}\right).$$

Then, for u > 0, setting $s = \{2\rho_j\sqrt{n} + Cc_{\mathcal{G}}/h\}\sqrt{u}$ yields

$$\mathbb{P}\left(\sup_{\theta\in\mathcal{G}_e(\rho_j)}|\mathbb{P}_n\theta|\leqslant 2\mathbb{E}\sup_{\theta\in\mathcal{G}_e(\rho_j)}|\mathbb{P}_n\theta|+\frac{2}{\sqrt{n}}\rho_j\sqrt{u}+C\frac{c_{\mathcal{G}}}{hn}\sqrt{u}\right)\geqslant 1-4e^{-u}.$$
 (B.21)

Then, conforming to the notation in the proof of Theorem 4 in Giné et al. (2003), and essentially by setting $u = s + 2\log(l - j + 1)$ in the event in the probability above, we define the events

$$\mathcal{E}_{j}^{+}(s) = \Big\{ \sup_{\theta \in \mathcal{G}_{e}(\rho_{j})} |\mathbb{P}_{n}\theta| \leq 2\mathbb{E} \sup_{\theta \in \mathcal{G}_{e}(\rho_{j})} |\mathbb{P}_{n}\theta| \\ + \frac{2\sqrt{s+2\log(l-j+1)}}{\sqrt{n}}\rho_{j} + C\frac{c_{\mathcal{G}}}{hn}\sqrt{s+2\log(l-j+1)} \Big\}.$$

Then $\mathbb{P}(\bigcap_{j=1}^{l} \mathcal{E}_{j}^{+}(s)) \ge 1-8e^{-s}$. Further introduce the notation that $\omega_{n}(a) = n^{1/2}\mathbb{E}\sup_{\theta\in\mathcal{G}_{e}(\rho_{j})} |(\mathbb{P}_{n}-P)\theta|$ if $a \in (\rho_{j-1}, \rho_{j}]$. We now find a function $\omega : [0, s_{n}] \to \mathbb{R}^{+}$ that satisfies the conditions in Theorem 4 in Giné et al. (2003), that is, (i) ω is non-decreasing; (ii) $\omega(a)/a$ is non-increasing; (iii) $\omega_{n}(a) \le \omega(a), \forall a \in [r_{n}, s_{n}]$; (iv) $\sup_{a \in (0,1]} a \sqrt{\log \log(1/a)}/\omega(a) \le K < \infty$. We set out to find such a ω .

Similar to the proof of Proposition B.7, let $N(\mathcal{G}_e(\rho_j), L_2(Q), \tau)$ be the covering number for the class $\mathcal{G}_e(\rho_j)$ by balls of radius τ in the $L_2(Q)$ norm for a generic measure Q on $(Y, \mathbf{F}, \mathbf{X})$, and let the constants A_1, A_2 be as in Theorem 3.5.4 in Giné and Nickl (2016). Then, by that theorem, if $a \in (\rho_{j-1}, \rho_j]$,

$$\omega_{n}(a) \leq A_{1} \frac{c_{\mathcal{G}}}{h} \int_{0}^{\rho_{j}/\|G\|_{L_{2}(Q)}} \sup_{Q} \sqrt{\log\{2N(\mathcal{G}_{e}(\rho_{j}), L_{2}(Q), \tau \|G\|_{L_{2}(Q)})\}} d\tau
+ A_{2} \frac{c_{\mathcal{G}}}{h} \left\{ \int_{0}^{\rho_{j}/\|G\|_{L_{2}(Q)}} \sup_{Q} \sqrt{\log\{2N(\mathcal{G}_{e}(\rho_{j}), L_{2}(Q), \tau \|G\|_{L_{2}(Q)})\}} d\tau \right\}^{2} / \sqrt{n}\rho_{j}^{2}.$$

By Eq. (A.2), $N(\mathcal{G}_n, L_2(Q), \tau) \leq N(\mathcal{F}_n(\overline{r}+1), L_2(Q), c_-h\tau)$ for a small enough constant $c_- > 0$ that is not dependent on Q. Then, by the derivation leading to Eq. (B.17), for Q the set of all generic measures Q on (Y, F, X),

$$\sup_{Q \in \mathcal{Q}} \log\{2N(\mathcal{G}_{e}(\rho_{j}), L_{2}(Q), \tau \|G\|_{L_{2}(Q)})\} \leq C_{1} V_{\mathcal{F}_{n}(\bar{r}+1)} \log(1/c_{-}h\tau).$$
(B.22)

Thus, for a large enough constant C_2 , we can set

$$\omega(u) = C_2 \left\{ V_{\mathcal{F}_n(\bar{r}+1)}^{1/2} \log^{1/2} \left(\frac{1}{h^2 u} \right) u + \frac{1}{n^{1/2} h} V_{\mathcal{F}_n(\bar{r}+1)} \log \left(\frac{1}{h^2 u} \right) \right\}.$$

Now, on the event $\cap_{j=1}^{l} \mathcal{E}_{j}^{+}(s), \forall j, \forall \theta \in \mathcal{G}_{e}(\rho_{j})$, by condition (B.19),

$$\sqrt{n} |\mathbb{P}_n \theta| \leq 2\omega_n (\|\theta\|_{L_2}) + 2\sqrt{s + 2\log(l - j + 1)}\rho_j + C\frac{c_{\mathcal{G}}}{h\sqrt{n}}\sqrt{s + 2\log(l - j + 1)} \\ \leq 2\omega(\rho_j) + (2 + Cc_{\mathcal{G}})\sqrt{s + 2\log(l - j + 1)}\rho_j.$$

From condition B.19, we have $1/(h\sqrt{n}) \lesssim r_n \lesssim \rho_j$. By our condition (iv) on ω and the derivations in the proof of Theorem 4 in Giné et al. (2003),

$$\max_{1 \leq j \leq l} \frac{\rho_j \sqrt{\log(l-j+1)}}{\omega(\rho_j)} \leq C_3.$$

Hence, on the same event (that is, the event $\bigcap_{j=1}^{l} \mathcal{E}_{j}^{+}(s)$), $\forall j, \forall \theta \in \mathcal{G}_{e}(\rho_{j})$,

$$\begin{aligned} |\mathbb{P}_n\theta| &\leq \{2 + (2 + Cc_{\mathcal{G}})\sqrt{2}C_3\}\frac{1}{\sqrt{n}}\omega(\rho_j) + (2 + Cc_{\mathcal{G}})\rho_j\sqrt{\frac{s}{n}} \\ &\leq C\left\{\left(\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\right)\|\theta\|_{L_2} + \frac{1}{h}\nu_{\mathrm{s},n}^2\right\}. \end{aligned}$$

In addition, for the class $\mathcal{G}(r_n)$, we can just invoke (B.21) but with $\mathcal{G}_e(\rho_j)$ replaced by $\mathcal{G}(r_n)$, and simplify with (B.19), to conclude that

$$\mathbb{P}\left[\sup_{\theta\in\mathcal{G}(r_n)}\sqrt{n}\left|\mathbb{P}_n\theta\right|\leqslant C\left\{\frac{1}{h}\nu_{\mathrm{s},n}^2+\nu_{\mathrm{s},n}\sqrt{\frac{s}{n}}\right\}\right]\geqslant 1-4e^{-s}$$

Combining the two cases above, we obtain the conclusion of the proposition.

C Proofs of main theorems and additional results for Section 3

C.1 Proof of Theorem 3.1

Proof. We first introduce a shorthand notation that we will employ throughout: for a possibly random function m = m(f, x),

$$\mathbb{P}m = \int m(\boldsymbol{f}, \boldsymbol{x}) \mathrm{d}\mathbb{P} = \int m(\boldsymbol{f}, \boldsymbol{x}) \mathrm{d}\mathbb{P}.$$

Since $\widetilde{F} = \widetilde{F}(X)$, for a possibly random function $g = g(\widetilde{f}, x_j)$, the analogous expression

$$\mathbb{P}g = \int g(\widetilde{f}, x_j) d\mathbb{P} = \int g(\widetilde{f}, x_j) d\mathbb{P}$$

makes sense as well. We note that $\mathbb{P}m$ and $\mathbb{P}g$ could still be random due to the potential inherent randomness of m and g.

We first prove (22). Recall \tilde{g}_n from Section B.2. To start, we have

$$\mathbb{P}(\widehat{g}_n - m_0)^2 = \mathbb{P}(\widehat{g}_n - m_0)^2 - \mathbb{P}_n(\widehat{g}_n - m_0)^2
+ \mathbb{P}_n(\widehat{g}_n - m_0)^2 - \mathbb{P}_n\ell(\cdot;\widehat{g}_n) + \mathbb{P}_n\ell(\cdot;\widehat{g}_n) - \mathbb{P}_n\ell(\cdot;\widetilde{g}_n) + \mathbb{P}_n\ell(\cdot;\widetilde{g}_n)
\leq \mathbb{P}(\widehat{g}_n - m_0)^2 - \mathbb{P}_n(\widehat{g}_n - m_0)^2 + \left\{\mathbb{P}_n(\widehat{g}_n - m_0)^2 - \mathbb{P}_n\ell(\cdot;\widehat{g}_n)\right\} + \mathbb{P}_n\ell(\cdot;\widetilde{g}_n), \quad (C.1)$$

where the inequality follows by the definition of \widehat{g}_n (in particular that it minimizes the loss $\mathbb{P}_n\ell(\cdot;g)$ within $g \in \mathcal{F}_n(\overline{r}+1)$). Next, for the term in the curly bracket in the last step of (C.1),

$$\mathbb{P}_{n}(\widehat{g}_{n} - m_{0})^{2} - \mathbb{P}_{n}\ell(\cdot;\widehat{g}_{n}) = \mathbb{P}_{n}(\widehat{g}_{n} - m_{0})^{2} - \mathbb{P}_{n}(\widehat{g}_{n} - m_{0})^{2} + \frac{2}{n}\sum_{i=1}^{n}\epsilon_{i}(\widehat{g}_{n}(\widetilde{F}_{i}, X_{i,j}) - m_{0}(F_{i}, X_{i,j})) - \frac{1}{n}\sum_{i=1}^{n}\epsilon_{i}^{2} = \frac{2}{n}\sum_{i=1}^{n}\epsilon_{i}(\widehat{g}_{n}(\widetilde{F}_{i}, X_{i,j}) - m_{0}(F_{i}, X_{i,j})) - \frac{1}{n}\sum_{i=1}^{n}\epsilon_{i}^{2}$$

and, for the last term in the last step of (C.1),

$$\mathbb{P}_n\ell(\cdot;\widetilde{g}_n) = \mathbb{P}_n(\widetilde{g}_n - m_0)^2 - \frac{2}{n}\sum_{i=1}^n \epsilon_i(\widetilde{g}_n - m_0) + \frac{1}{n}\sum_{i=1}^n \epsilon_i^2.$$

Plugging the two equalities above into (C.1),

$$\mathbb{P}(\hat{g}_{n} - m_{0})^{2} \leq \underbrace{\mathbb{P}(\hat{g}_{n} - m_{0})^{2} - \mathbb{P}_{n}(\hat{g}_{n} - m_{0})^{2}}_{I_{n,1}} \\
+ \frac{2}{n} \sum_{i=1}^{n} \epsilon_{i}(\hat{g}_{n}(\widetilde{F}_{i}, X_{i,j}) - m_{0}(F_{i}, X_{i,j})) - \frac{2}{n} \sum_{i=1}^{n} \epsilon_{i}(\widetilde{g}_{n} - m_{0}) + \mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2} \\
= I_{n,1} + \frac{2}{n} \sum_{i=1}^{n} \epsilon_{i}(\hat{g}_{n} - \widetilde{g}_{n})(\widetilde{F}_{i}, X_{i,j}) + \mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2}.$$
(C.2)

We next analyze $I_{n,1}$. Note that the class of functions $\{g - m_0 = g(\tilde{f}, x_j) - m_0(f, x_j) : g \in \mathcal{F}_n(\bar{r}+1)\}$ satisfies the covering number bound in Lemma B.6. By Lemma B.6, on an event $\mathcal{A}_{1,\zeta,s'}$ satisfying (B.12), we have

$$\forall g \in \mathcal{F}_n(\overline{r}+1), \quad |\mathbb{P}(g-m_0)^2 - \mathbb{P}_n(g-m_0)^2| \le \zeta \left\{ \mathbb{P}(g-m_0)^2 + \frac{s'}{n} \right\}.$$
 (C.3)

Applying this result with $g = \hat{g}_n$ on (C.2), on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'}$ we obtain

$$\mathbb{P}(\widehat{g}_n - m_0)^2 \leqslant \frac{s'}{n} + 4 \underbrace{\frac{1}{n} \sum_{i=1}^n \epsilon_i(\widehat{g}_n - \widetilde{g}_n)(\widetilde{F}_i, X_{i,j})}_{I_{n,2}} + 2\mathbb{P}_n(\widetilde{g}_n - m_0)^2.$$
(C.4)

Next, we investigate the rate of the stochastic term $I_{n,2}$ on the right-hand side of (C.4). Recalling the $\tilde{\epsilon}_i$'s introduced in Section B.4, we further decompose

$$I_{n,2} = \frac{1}{n} \sum_{i=1}^{n} \widetilde{\epsilon}_i (\widehat{g}_n - \widetilde{g}_n) (\widetilde{F}_i, X_{i,j}) + \frac{1}{n} \sum_{i=1}^{n} (\epsilon_i - \widetilde{\epsilon}_i) (\widehat{g}_n - \widetilde{g}_n) (\widetilde{F}_i, X_{i,j}) \equiv I_{n,2,1} + I_{n,2,2}.$$

The reason we need this decomposition is that in $I_{n,2,1}$, the $\tilde{\epsilon}_i$'s are conditionally properly centered, that is $\mathbb{E}[\tilde{\epsilon}_i | \tilde{F}_i, X_{i,j}] = 0$, while the term $I_{n,2,2}$ will utilize the "stability of conditional expectation" condition in (B.1) to treat the difference $\epsilon_i - \tilde{\epsilon}_i$.

Now, for the term $I_{n,2,1}$, note that the class of functions $\{(g - \tilde{g}_n) = (g - \tilde{g}_n)(\tilde{f}, x_j) : g \in \mathcal{F}_n(\bar{r}+1)\}$ satisfies the covering number bound in Lemma B.6; we then apply (B.14) in Proposition B.7 with $\theta = \hat{g}_n - \tilde{g}_n$ to conclude that, on an event $\mathcal{A}_{2,s}$ satisfying $\mathbb{P}(\mathcal{A}_{2,s}) \geq 1 - 12e^{-s}$,

$$|I_{n,2,1}| = \left|\frac{1}{n}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}(\widehat{g}_{n} - \widetilde{g}_{n})(\widetilde{F}_{i}, X_{i,j})\right| \leq c_{4}\left\{\|\widehat{g}_{n} - \widetilde{g}_{n}\|_{L_{2}} + \nu_{\mathrm{s},n}\right\}\left\{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\right\}.$$
 (C.5)

For the term $I_{n,2,2}$, first note that $|\epsilon_i - \tilde{\epsilon}_i| = |\{Y_i - m_0(\mathbf{F}_i, X_{i,j})\} - \{Y_i - \mathbb{E}[Y_i|\tilde{\mathbf{F}}_i, X_{i,j}]\}| = |m_0(\mathbf{F}_i, X_{i,j}) - \mathbb{E}[Y|\tilde{\mathbf{F}}, X_j]|$, which is uniformly bounded by $2M_\infty$ by the comments in the proof or Proposition B.7. As already mentioned, the class of functions $\{(g - \tilde{g}_n) = (g - \tilde{g}_n)(\tilde{\mathbf{f}}, x_j) : g \in \mathcal{F}_n(\bar{r}+1)\}$ satisfies the covering number bound in Lemma B.6; multiply this class by $\epsilon - \tilde{\epsilon} = \mathbb{E}[Y|\tilde{\mathbf{F}}, X_j] - m_0(\mathbf{F}_i, X_{i,j})$ considered as a single, uniformly bounded function changes at most the constant in the covering number bound (see Theorem 3 in Andrews (1994)). Hence, we can apply (B.15) in Proposition B.7 to conclude that, on an event $\mathcal{A}_{3,s}$ satisfying $\mathbb{P}(\mathcal{A}_{3,s}) \geq 1 - 12e^{-s}$,

$$|I_{n,2,2} - \int I_{n,2,2} d\mathbb{P}| \leq c_4 \left\{ \| (\epsilon - \widetilde{\epsilon}) (\widehat{g}_n - \widetilde{g}_n) (\widetilde{F}, X_j) \|_{L_2} + \nu_{\mathrm{s},n} \right\} \left\{ \nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}} \right\}$$
$$\lesssim \left\{ \| \widehat{g}_n - \widetilde{g}_n \|_{L_2} + \nu_{\mathrm{s},n} \right\} \left\{ \nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}} \right\}$$

where in the transition to the last line we have again applied the boundedness of $\epsilon - \tilde{\epsilon}$. For the centering term $\int I_{n,2,2} d\mathbb{P}$, by the Cauchy-Schwarz inequality,

$$\left|\int I_{n,2,2} \mathrm{d}\mathbb{P}\right| \leq \left\{\mathbb{E}\left[\left(\epsilon_{i}-\widetilde{\epsilon}_{i}\right)^{2}\right]\right\}^{1/2} \left\{\mathbb{P}(\widehat{g}_{n}-\widetilde{g}_{n})^{2}\right\}^{1/2} \lesssim \delta_{\mathrm{f}} \|\widehat{g}_{n}-\widetilde{g}_{n}\|_{L_{2}}.$$
(C.6)

Collecting the rates for $I_{n,2,1}$ and $I_{n,2,2}$, we conclude that on $\mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$,

$$|I_{n,2}| = |I_{n,2,1} + I_{n,2,2}| \lesssim \{\|\widehat{g}_n - \widetilde{g}_n\|_{L_2} + \nu_{s,n}\} \{\nu_{s,n} + \sqrt{\frac{s}{n}}\} + \delta_f \|\widehat{g}_n - \widetilde{g}_n\|_{L_2}.$$

Then, applying this result in (C.4), we obtain that on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$ we have (note the change in notation from $\mathbb{P}(\widehat{g}_n - m_0)^2$ to $\|\widehat{g}_n - m_0\|_{L_2}^2$ below)

$$\begin{aligned} \|\widehat{g}_{n} - m_{0}\|_{L_{2}}^{2} &\leq \frac{s'}{n} + C \left\{ \|\widehat{g}_{n} - \widetilde{g}_{n}\|_{L_{2}} + \nu_{\mathrm{s},n} \right\} \left\{ \nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}} \right\} + C\delta_{\mathrm{f}} \|\widehat{g}_{n} - \widetilde{g}_{n}\|_{L_{2}} + 2\mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2} \\ &\lesssim \|\widehat{g}_{n} - \widetilde{g}_{n}\|_{L_{2}} \{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \nu_{\mathrm{s},n} \{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\} + \frac{s'}{n} + 2\mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2} \end{aligned}$$

$$\lesssim \{\|\widehat{g}_{n} - m_{0}\|_{L_{2}} + \|\widetilde{g}_{n} - m_{0}\|_{L_{2}}\}\{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \nu_{\mathrm{s},n}\{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\} + \frac{s'}{n} + 2\mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2} \\ \lesssim \|\widehat{g}_{n} - m_{0}\|_{L_{2}}\{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \{\nu_{\mathrm{a},n} + \delta_{\mathrm{f}}\}\{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \nu_{\mathrm{s},n}\{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\} + \frac{s'}{n} \\ + 2\underbrace{\mathbb{P}_{n}(\widetilde{g}_{n} - m_{0})^{2}}_{I_{n,3}}$$
(C.7)

where in the transition to the last step we have invoked Lemma B.1. For the term $I_{n,3}$, by (C.3) and the remark above it, and then Lemma B.1 again, on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'}$ we have

$$I_{n,3} = \mathbb{P}_n (\tilde{g}_n - m_0)^2 \leq \frac{3}{2} \mathbb{P} (\tilde{g}_n - m_0)^2 + \frac{1}{2} \frac{s'}{n} \lesssim \nu_{\mathrm{a},n}^2 + \delta_{\mathrm{f}}^2 + \frac{s'}{n}.$$

Then, invoking the above in (C.7) and then using $ab \leq 2^{-1}(a^2+b^2)$ several times, we obtain that on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$,

$$\begin{split} \|\widehat{g}_{n} - m_{0}\|_{L_{2}}^{2} &\lesssim \|\widehat{g}_{n} - m_{0}\|_{L_{2}} \{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} \\ &+ \{\nu_{\mathrm{a},n} + \delta_{\mathrm{f}}\}\{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \nu_{\mathrm{s},n}\{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\} + \nu_{\mathrm{a},n}^{2} + \delta_{\mathrm{f}}^{2} + \frac{s'}{n} \\ &\lesssim \|\widehat{g}_{n} - m_{0}\|_{L_{2}}\{\nu_{\mathrm{s},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}}\} + \{\nu_{\mathrm{s},n}^{2} + \nu_{\mathrm{a},n}^{2} + \delta_{\mathrm{f}}^{2} + \frac{s}{n} + \frac{s'}{n}\}. \end{split}$$

Finally, solving for $\|\widehat{g}_n - m_0\|_{L_2}$ yields that, on $\mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$,

$$\|\hat{g}_n - m_0\|_{L_2} \lesssim \nu_{\mathrm{s},n} + \nu_{\mathrm{a},n} + \delta_{\mathrm{f}} + \sqrt{\frac{s}{n}} + \sqrt{\frac{s'}{n}}.$$
 (C.8)

Next, we impose Assumption 3.2 and derive the concrete resultant rates for $\nu_{s,n}$ in (B.11) (which then becomes the rate ν_n in (21)) and $\nu_{a,n}$. For the former, namely $\nu_{s,n}$, essentially we need to obtain an upper bound for the VC-index $V_{\mathcal{F}_n(\bar{r}+1)}$ of the class $\mathcal{F}_n(\bar{r}+1)$; see our Def. A.3. Note that there are a couple of closely-related though not always identical concepts regarding the VC-index. Our definition turns out to be identical to the *pseudo*dimension defined in, for instance, Def. 11.2 in Anthony and Bartlett (1999) and Def. 2 in Bartlett et al. (2019) for real-valued functions. A closely related concept, termed VCdim(\mathcal{G}) for a class \mathcal{G} of $\{0,1\}$ -valued functions, also appears; for instance, see Theorem 14.1 in Anthony and Bartlett (1999). Fortunately, we have the following simple connection: let \mathcal{F}_n^{\dagger} be our class of function $\mathcal{F}_n(\bar{r}+1)$ augmented with one extra input unit and one extra computation unit, and let \mathcal{H}^{\dagger} be the resulting set of $\{0, 1\}$ -valued functions computed by \mathcal{F}_n^{\dagger} , all as described in Theorem 14.1 in Anthony and Bartlett (1999). Then the VC-index of $\mathcal{F}_n(\bar{r}+1)$ will be bounded by $\operatorname{VCdim}(\mathcal{H}^{\dagger})$, that is $V_{\mathcal{F}_n(\bar{r}+1)} \leq \operatorname{VCdim}(\mathcal{H}^{\dagger})$, and the right-hand side is in turn bounded in Anthony and Bartlett (1999). Specifically, let U be the total number of weights for a function in $\mathcal{F}_n(\bar{r}+1)$, so following the derivation on p. 2 in appendix B of Kohler and Langer (2021), $U \sim Lk_0^2$ (note that the input dimension \overline{r} + 1 does not appear because by assumption \overline{r} is a constant). Then, by Theorem 7 in Bartlett et al. (2019) and then Assumption 3.2, $V_{\mathcal{F}_n(\bar{r}+1)} \leq \operatorname{VCdim}(\mathcal{H}^{\dagger}) \leq UL\log(U) \sim$ $n^{\frac{1}{2\kappa+1}}\log^{2\frac{4\kappa-1}{2\kappa+1}+1}(n)$. To fix the multiplicative factor in this bound on $V_{\mathcal{F}_n(\bar{r}+1)}$, we shall explicitly set

$$V_{\mathcal{F}_n(\bar{r}+1)} \leqslant c_p n^{\frac{1}{2\kappa+1}} \log^{2\frac{4\kappa-1}{2\kappa+1}+1}(n) = p_n \log^{-1}(n)$$
(C.9)

for a large enough constant c_p and for p_n from (21).

Now we are ready to fix s and s' in (C.8). We set $s = p_n$ so $\sqrt{s/n}$ precisely becomes the concrete rate ν_n in (21) (up to a factor involving c_p), and ν_n in turn bounds $\nu_{s,n}$ from (B.11) if we apply in (B.11) the bound (C.9). Similarly, we choose s' to be a large enough multiple of p_n . Together with (C.8), we then conclude that, on the intersection $\mathcal{A}' \equiv \mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$ whose probability is at least $1 - Ce^{-p_n}$,

$$\|\hat{g}_n - m_0\|_{L_2} \lesssim \nu_n + \nu_{\mathrm{a},n} + \delta_{\mathrm{f}}.$$
 (C.10)

Finally, we derive in (C.10) the rate of $\nu_{a,n}$ in order to reach (22). This rate is implicitly implied in Corollary 1 in Fan and Gu (2023), and depends on the earlier Proposition 3.3 in Fan et al. (2022), but we give a brief derivation here. Define \tilde{L} and \tilde{k}_0 such that $\tilde{L}\log(\tilde{L}) = L$ and $\tilde{k}_0\log(\tilde{k}_0) = k_0$; we introduce these quantities because the neural network architecture scaling in Proposition 3.3 in Fan et al. (2022) is specified by such quantities. Under Assumption 3.2, it's easy to compute that $\tilde{L}\tilde{k}_0 \sim n^{\frac{1}{4\kappa+2}}\log^{\frac{4\kappa-1}{2\kappa+1}-2}(n)$, which further yields, by Proposition 3.3 in Fan et al. (2022),

$$\nu_{\mathbf{a},n} \lesssim (\widetilde{L}\widetilde{k}_0)^{-2\kappa} \sim \nu_n. \tag{C.11}$$

Finally, applying the bound above in (C.10) yields that (22) holds on the intersection $\mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s} \cap \mathcal{A}_{3,s}$.

Finally, given (22), a straightforward application of Lemma A.4 with $\hat{g} = \hat{g}_n$ and $m = m_0$ will yield (23). This completes the proof of the theorem.

C.2 Proof of Theorem 3.2

Proof. Introduce the shorthand notation $\mathbb{P}_{n|\mathcal{I}_h}$ for the empirical measure \mathbb{P}_n but only including the samples with sample indices in \mathcal{I}_h . Then, the loss function $\hat{R}_n^{\text{null}}(\cdot)$ introduced in Section 2.4 can be written as

$$\widehat{R}_n^{\text{null}}(\alpha) = \mathbb{P}_n \alpha^2 - 2\mathbb{P}_{n|\mathcal{I}_h} \alpha_j^{\text{s}}.$$

Also recall $\tilde{\alpha}_n$ defined in (B.3). Then, we have

$$\mathbb{P}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} = \mathbb{P}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} - \mathbb{P}_{n}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} + \mathbb{P}_{n}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} \\
- \widehat{R}_{n}^{\text{null}}(\widehat{\alpha}_{n}) + \widehat{R}_{n}^{\text{null}}(\widehat{\alpha}_{n}) - \widehat{R}_{n}^{\text{null}}(\widetilde{\alpha}_{n}) + \widehat{R}_{n}^{\text{null}}(\widetilde{\alpha}_{n}) - \mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} + \mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} \\
\leq \mathbb{P}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} - \mathbb{P}_{n}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} + \underbrace{\mathbb{P}_{n}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} - \widehat{R}_{n}^{\text{null}}(\widehat{\alpha}_{n})}_{I_{n,2}} \\
+ \underbrace{\widehat{R}_{n}^{\text{null}}(\widetilde{\alpha}_{n}) - \mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2}}_{I_{n,3}} + \mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2},$$
(C.12)

where the inequality is due to $\hat{\alpha}_n$ minimizing the loss $\hat{R}_n^{\text{null}}(\cdot)$. Now, for handling the terms $I_{n,2}$ and $I_{n,3}$ in the above, note that for any $\alpha \in \mathcal{F}_n$,

$$\mathbb{P}_n(\alpha - \alpha_n^{\text{null}})^2 - \widehat{R}_n^{\text{null}}(\alpha) = \mathbb{P}_n \alpha^2 - 2\mathbb{P}_n(\alpha \alpha_n^{\text{null}}) + \mathbb{P}_n \alpha_n^{2\text{null}} + 2\mathbb{P}_{n|\mathcal{I}_h} \alpha_j^{\text{s}} - \mathbb{P}_n \alpha^2$$
$$= -2\mathbb{P}_n(\alpha \alpha_n^{\text{null}}) + \mathbb{P}_n \alpha_n^{2\text{null}} + 2\mathbb{P}_{n|\mathcal{I}_h} \alpha_j^{\text{s}}.$$

Then, together using the function $\theta(\cdot; \alpha)$ introduced in (B.18), the sum of the terms $I_{n,2}$ and $I_{n,3}$ becomes

$$\mathbb{P}_{n}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} - \widehat{R}_{n}^{\text{null}}(\widehat{\alpha}_{n}) + \widehat{R}_{n}^{\text{null}}(\widetilde{\alpha}_{n}) - \mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} \\
= -2\mathbb{P}_{n}(\widehat{\alpha}_{n}\alpha_{n}^{\text{null}}) + 2\mathbb{P}_{n}(\widetilde{\alpha}_{n}\alpha_{n}^{\text{null}}) + 2\mathbb{P}_{n|\mathcal{I}_{h}}\widehat{\alpha}_{n,j}^{\text{s}} - 2\mathbb{P}_{n|\mathcal{I}_{h}}\widetilde{\alpha}_{n,j}^{\text{s}} = 2\mathbb{P}_{n}\theta(\cdot;\widehat{\alpha}_{n}). \quad (C.13)$$

As already mentioned in Section B.5, The function θ captures the deviation from the null, which we will explain later.

Plugging (C.13) into (C.12), we obtain

$$\mathbb{P}(\widehat{\alpha}_n - \alpha_n^{\text{null}})^2 \leq \underbrace{\mathbb{P}(\widehat{\alpha}_n - \alpha_n^{\text{null}})^2 - \mathbb{P}_n(\widehat{\alpha}_n - \alpha_n^{\text{null}})^2}_{I_{n,1}} + 2\mathbb{P}_n\theta(\cdot;\widehat{\alpha}_n) + \mathbb{P}_n(\widetilde{\alpha}_n - \alpha_n^{\text{null}})^2. \quad (C.14)$$

For the term $I_{n,1}$ above, first note that the class of functions $\{\alpha - \alpha_n^{\text{null}} : \alpha \in \mathcal{F}_n(\overline{r} + 1)\}$ satisfies the covering number bound (A.1) by our assumptions. Then, similar to (C.3), by Lemma B.6, on an event $\mathcal{A}_{1,\zeta,s'}$ satisfying (B.12), we have

$$\forall \alpha \in \mathcal{F}_n(\overline{r}+1), \quad |\mathbb{P}(\alpha - \alpha_n^{\text{null}})^2 - \mathbb{P}_n(\alpha - \alpha_n^{\text{null}})^2| \leq \zeta \left\{ \mathbb{P}(\alpha - \alpha_n^{\text{null}})^2 + \frac{s'}{n} \right\}.$$
(C.15)

Applying this result with $\alpha = \hat{\alpha}_n$ for $I_{n,1}$ on (C.14), on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'}$ we obtain

$$\mathbb{P}(\widehat{\alpha}_n - \alpha_n^{\text{null}})^2 \leqslant \frac{1}{2} \mathbb{P}(\widehat{\alpha}_n - \alpha_n^{\text{null}})^2 + \frac{1}{2}\frac{s'}{n} + 2\mathbb{P}_n\theta(\cdot;\widehat{\alpha}_n) + \mathbb{P}_n(\widetilde{\alpha}_n - \alpha_n^{\text{null}})^2$$

or equivalently, rearranging and centering we have,

$$\mathbb{P}(\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} \leqslant \frac{s'}{n} + 4\mathbb{P}_{n}\theta(\cdot;\widehat{\alpha}_{n}) + 2\mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2} \\
= \frac{s'}{n} + 4\underbrace{\mathbb{P}_{n}\theta(\cdot;\widehat{\alpha}_{n})}_{I_{n,4}} + 2\underbrace{\mathbb{P}_{n}(\widetilde{\alpha}_{n} - \alpha_{n}^{\text{null}})^{2}}_{I_{n,5}}.$$
(C.16)

Next, we consider the term $I_{n,4}$ in (C.16). First, we note that

$$\begin{aligned} \|\theta(\cdot;\alpha)\|_{L_{2}}^{2} &= \mathbb{E}\left[\left\{(\alpha - \widetilde{\alpha}_{n})_{j}^{s}(\widetilde{\boldsymbol{F}}, X_{j})\mathbb{1}\{X_{j} \in \mathcal{B}_{h}\} - (\alpha - \widetilde{\alpha}_{n})\alpha_{n}^{\mathrm{null}}\}(\widetilde{\boldsymbol{F}}, X_{j})\right\}^{2}\right] \\ &\leq 2\mathbb{E}\left[(\alpha - \widetilde{\alpha}_{n})_{j}^{s}(\widetilde{\boldsymbol{F}}, X_{j})^{2}\mathbb{1}\{X_{j} \in \mathcal{B}_{h}\} + \{(\alpha - \widetilde{\alpha}_{n})\alpha_{n}^{\mathrm{null}}\}^{2}(\widetilde{\boldsymbol{F}}, X_{j})\right] \lesssim h^{-2}\|\alpha - \widetilde{\alpha}_{n}\|_{L_{2}}^{2}\end{aligned}$$

where the last inequality follows from (A.2). Then, applying Proposition B.8 (whose condition (B.19) is satisfied due to Assumptions 3.2, 3.5 and the paragraph below (C.8)) and the result above on the term $(\mathbb{P}_n - \mathbb{P})\theta(\cdot; \hat{\alpha}_n)$ in (C.16), we obtain that, on an event $\mathcal{A}_{2,s}$ satisfying $\mathbb{P}(\mathcal{A}_{2,s}) \ge 1 - 12e^{-s}$,

$$|I_{n,4}| \leq c_4 \left\{ \|\theta(\cdot; \hat{\alpha}_n)\|_{L_2} + \nu_{s,n} \right\} \left\{ \nu_{s,n} + \sqrt{\frac{s}{n}} \right\} \lesssim \left\{ \frac{1}{h} \|\hat{\alpha}_n - \alpha_n^{\text{null}}\|_{L_2} + \nu_{s,n} \right\} \left\{ \nu_{s,n} + \sqrt{\frac{s}{n}} \right\} + \frac{1}{h} \nu_{s,n}^2.$$

Next, we consider the term $I_{n,5}$ in (C.16), which represents a population bias induced by approximating α_n^{null} through the neural network function $\tilde{\alpha}_n$. By (C.15) and the remarks above, on the event $\mathcal{A}_{1,\zeta=\frac{1}{2},s'}$ we have

$$|I_{n,5}| = |\mathbb{P}_n(\widetilde{\alpha}_n - \alpha_n^{\text{null}})^2| \leq \frac{1}{2}\mathbb{P}(\widetilde{\alpha}_n - \alpha_n^{\text{null}})^2 + \frac{1}{2}\frac{s'}{n} \lesssim \nu_{\mathrm{a},n}^2 + \delta_{\mathrm{f}}^2 + \frac{s'}{n}$$

where the last step follows by Lemma B.4. Then, collecting the rates for $I_{n,4}$ and $I_{n,5}$ in (C.16), we obtain, on the intersection $\mathcal{A}_{1,\zeta=\frac{1}{2},s'} \cap \mathcal{A}_{2,s}$,

$$\|\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}}\|_{L_{2}}^{2} \lesssim h^{-1}(\nu_{\text{s},n} + \sqrt{\frac{s}{n}})\|\widehat{\alpha}_{n} - \alpha_{n}^{\text{null}}\|_{L_{2}} + \nu_{\text{s},n}\left\{\nu_{\text{s},n} + \sqrt{\frac{s}{n}}\right\} + \frac{1}{h}\nu_{\text{s},n}^{2} + \nu_{\text{a},n}^{2} + \delta_{\text{f}}^{2} + \frac{s'}{n}.$$

Then, solving for $\|\widehat{\alpha}_n - \alpha_n^{\text{null}}\|_{L_2}$ yields that

$$\|\widehat{\alpha}_n - \alpha_n^{\text{null}}\|_{L_2} \lesssim h^{-1}(\nu_{\mathbf{s},n} + \sqrt{\frac{s}{n}}) + \nu_{\mathbf{a},n} + \delta_{\mathbf{f}} + \sqrt{\frac{s'}{n}}.$$

This is our counterpart to (C.8) in the current proof of Theorem 3.2. Then, we apply Assumption 3.2 as in the paragraph below (C.8), and subsequently choose s and s', and bound $\nu_{s,n}$, as in the paragraph below (C.9), and bound $\nu_{a,n}$ as in (C.11). This will yield the rate, holding on an event with probability at least $1 - Ce^{-p_n}$, on the left-hand side of (27) except for the cap at M_{∞} . The additional cap is straightforward due the boundedness assumptions on α_n^{null} and M (which in turn truncates the neural network estimators) in Assumption 3.1. This completes the proof of Theorem 3.2 regarding (27).

Next, we bound $\|\alpha_n^{\text{null}} - \alpha_{t,n}^*\|_{L_2}$ which will subsequently lead to (29). We have, for all $t \in \mathbb{R}$,

$$\begin{aligned} \mathbb{P}(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})^{2} &= \mathbb{P}(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})^{2} - R_{n}^{\text{null}}(\alpha_{n}^{\text{null}}) + R_{n}^{\text{null}}(\alpha_{n}^{\text{null}}) - R_{n}^{\text{null}}(\alpha_{t,n}^{*}) + R_{n}^{\text{null}}(\alpha_{t,n}^{*}) \\ &\leq \mathbb{P}(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})^{2} - R_{n}^{\text{null}}(\alpha_{n}^{\text{null}}) + R_{n}^{\text{null}}(\alpha_{t,n}^{*}) \\ &= \|\alpha_{n}^{\text{null}}\|_{L_{2}}^{2} - 2\langle\alpha_{n}^{\text{null}}, \alpha_{t,n}^{*}\rangle + \|\alpha_{t,n}^{*}\|_{L_{2}}^{2} - \|\alpha_{n}^{\text{null}}\|_{L_{2}}^{2} + 2\langle\alpha_{n}^{\text{null}}, \alpha_{n}^{\text{null}}\rangle + \|\alpha_{t,n}^{*}\|_{L_{2}}^{2} - 2\langle\alpha_{t,n}^{*}, \alpha_{n}^{\text{null}}\rangle \\ &= -2\langle\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*}, \alpha_{t,n}^{*}\rangle + 2\langle\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*}, \alpha_{n}^{\text{null}}\rangle \\ &= -2\int_{\Omega_{h}} e^{t\Psi\{g_{0,j}^{s}(\tilde{f}, x_{j})\}}(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})_{j}^{s}(\tilde{f}, x_{j})d\mathbb{P} + 2\int_{\Omega_{h}}(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})_{j}^{s}(\tilde{f}, x_{j})d\mathbb{P} \\ &= -2\int_{\Omega_{h}} [e^{t\Psi\{g_{0,j}^{s}(\tilde{f}, x_{j})\}} - 1](\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})_{j}^{s}(\tilde{f}, x_{j})d\mathbb{P} \\ &\leq 2\|Z_{t,j,h}\|_{L_{2}}\|(\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*})_{j}^{s}\|_{L_{2}} \leqslant \frac{2c^{s}}{h}\|Z_{t,j,h}\|_{L_{2}}\|\alpha_{n}^{\text{null}} - \alpha_{t,n}^{*}\|_{L_{2}}, \end{aligned}$$
(C.17)

where the inequality is due to α_n^{null} minimizing the loss $R_n^{\text{null}}(\cdot)$, the transition to the second last line follows by the property of $\alpha_{t,n}^*$ as the Riesz representer in (25) and the corresponding property of α_n^{null} , the transition to the last line follows by the Cauchy-Schwarz inequality, and the last step follows by Lemma A.1. Then, combining (C.17) with the boundedness assumption on α_n^{null} and $\alpha_{t,n}^*$ yields (28).

Finally, applying the triangle inequality on (27) and (28) yields the remaining conclusion, namely (29), of Theorem 3.2.

C.3 Proof of Proposition 3.3

Proof. We first prove the first half of the proposition on $\hat{\delta}_t$. By definition,

$$\widehat{\delta}_{t} = \operatorname{argmin}_{\delta_{t} \in \mathbb{R}} \mathbb{P}_{n} \ell(\cdot; \widehat{g}_{n} + \delta_{t} \widehat{\alpha}_{n}).$$

Denote the function $\hat{\delta}_n = \hat{g}_n - m_0$. Then

$$\mathbb{P}_n\ell(\cdot;\hat{g}_n+\delta_{\mathsf{t}}\hat{\alpha}_n)=\frac{1}{n}\sum_{i=1}^n\left\{Y_i-\hat{g}_n(\widetilde{F}_i,X_{i,j})-\delta_{\mathsf{t}}\hat{\alpha}_n(\widetilde{F}_i,X_{i,j})\right\}^2$$

$$= \frac{1}{n} \sum_{i=1}^{n} \left\{ Y_i - \hat{g}_n(\widetilde{F}_i, X_{i,j}) \right\}^2 - 2\delta_t \frac{1}{n} \sum_{i=1}^{n} \left\{ Y_i - \hat{g}_n(\widetilde{F}_i, X_{i,j}) \right\} \hat{\alpha}_n(\widetilde{F}_i, X_{i,j}) + \delta_t^2 \frac{1}{n} \sum_{i=1}^{n} \hat{\alpha}_n^2(\widetilde{F}_i, X_{i,j}).$$

The last line above viewed as a quadratic function in δ_t reaches its minimum at $\delta_t = \hat{\delta}_t$ for $\hat{\delta}_t$ given in (15). We can also re-write (15) as

$$\widehat{\delta}_{t} = \frac{1}{\frac{1}{n}\sum_{i=1}^{n}\widehat{\alpha}_{n}^{2}(\widetilde{F}_{i}, X_{i,j})} \frac{1}{n} \sum_{i=1}^{n} \left\{ \epsilon_{i} - \widehat{\delta}_{n}(\widetilde{F}_{i}, X_{i,j}) \right\} \widehat{\alpha}_{n}(\widetilde{F}_{i}, X_{i,j})$$
(C.18)

$$= \underbrace{\frac{1}{\underbrace{\frac{1}{n}\sum_{i=1}^{n}\widehat{\alpha}_{n}^{2}(\widetilde{F}_{i}, X_{i,j})}_{J_{n,1}}}_{J_{n,1}}\underbrace{\frac{1}{n}\sum_{i=1}^{n}\epsilon_{i}\widehat{\alpha}_{n}(\widetilde{F}_{i}, X_{i,j})}_{J_{n,2}} - \underbrace{\frac{1}{\frac{1}{n}\sum_{i=1}^{n}\widehat{\alpha}_{n}^{2}(\widetilde{F}_{i}, X_{i,j})}_{J_{n,2}}}_{J_{n,3}}\underbrace{\frac{1}{n}\sum_{i=1}^{n}\widehat{\delta}_{n}(\widetilde{F}_{i}, X_{i,j})\widehat{\alpha}_{n}(\widetilde{F}_{i}, X_{i,j})}_{J_{n,3}}}_{J_{n,3}}$$

For the term $J_{n,1}$ (in the denominator) above, similar to (C.15) (with the choice s' being a large enough multiple of p_n), we conclude that on an event $\mathcal{A}_{1,1}$ with $\mathbb{P}(\mathcal{A}_{1,1}) \ge 1 - Ce^{-p_n}$,

$$\frac{1}{n}\sum_{i=1}^{n}\widehat{\alpha}_{n}^{2}(\widetilde{F}_{i},X_{i,j}) \geq \frac{1}{2}\int\widehat{\alpha}_{n}^{2}(\widetilde{f},x_{j})\mathrm{d}\mathbb{P} - C\nu_{n}^{2}.$$

Next,

$$\begin{split} &\int \widehat{\alpha}_{n}^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} \geq \int (\alpha_{n}^{\mathrm{null}})^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} + \int (\widehat{\alpha}_{n}^{2} - (\alpha_{n}^{\mathrm{null}})^{2})(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} \\ &= \int (\alpha_{n}^{\mathrm{null}})^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} + \int (\widehat{\alpha}_{n} + \alpha_{n}^{\mathrm{null}})(\widehat{\alpha}_{n} - \alpha_{n}^{\mathrm{null}})(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} \\ &\geq \int (\alpha_{n}^{\mathrm{null}})^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} - CM_{\infty} \{\int (\widehat{\alpha}_{n} - \alpha_{n}^{\mathrm{null}})^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} \}^{1/2} \geq \int (\alpha_{n}^{\mathrm{null}})^{2}(\widetilde{f}, x_{j}) \mathrm{d}\mathbb{P} - C(h^{-1}\nu_{n} + \delta_{\mathrm{f}}) \end{split}$$

where the last step holds on an event $\mathcal{A}_{1,2}$ with $\mathbb{P}(\mathcal{A}_{1,2}) \ge 1 - Ce^{-p_n}$ by Theorem 3.2. Combining the above two inequalities, by the conditions imposed in Proposition 3.3 with Assumption 3.6(i), on the event $\mathcal{A}_{1,1} \cap \mathcal{A}_{1,2}$ we have that $J_{n,1}$ is bounded from below from zero by a constant (independent of n).

Next, for the term $J_{n,2}$, we can follow the treatment for the terms $I_{n,2} = I_{n,2,1} + \{I_{n,2,2} - \int I_{n,2,2} d\mathbb{P}\} + \int I_{n,2,2} d\mathbb{P}$ in the proof of Theorem 3.1 in Section C.1, together with the choice of s in that proof, to conclude that on an event \mathcal{A}_2 with $\mathbb{P}(\mathcal{A}_2) \ge 1 - Ce^{-p_n}$,

$$|J_{n,2}| \lesssim \{\|\widehat{\alpha}_n\|_{L_2} + \nu_n\}\nu_n + \{\|\widehat{\alpha}_n\|_{L_2} + \nu_n\}\nu_n + \delta_{\mathrm{f}}\|\widehat{\alpha}_n\|_{L_2} \lesssim r_{m,n}$$

where the last step follows by the boundedness of $\hat{\alpha}_n$.

For the term $J_{n,3}$ above, first note that the class of functions $\{(g-m_0)\alpha : g, \alpha \in \mathcal{F}_n(\overline{r}+1)\}$ satisfies the covering number bound by Theorem 3 in Andrews (1994) and our assumptions. Then, by (B.15) in Proposition B.7, together with the choice of s in the proof of Theorem 3.1 in Section C.1, we conclude that an event \mathcal{A}_3 with $\mathbb{P}(\mathcal{A}_3) \ge 1 - Ce^{-p_n}$,

$$\begin{aligned} |J_{n,3} - \int J_{n,3} \mathrm{d}\mathbb{P}| &\lesssim \{ \|\widehat{\delta}_n \widehat{\alpha}_n\|_{L_2} + \nu_n \} \nu_n \lesssim \{ M \|\widehat{\delta}_n\|_{L_2} + \nu_n \} \nu_n \lesssim r_{m,n}, \\ |\int J_{n,3} \mathrm{d}\mathbb{P}| &\lesssim \{ (\mathbb{P}\widehat{\alpha}_n^2) (\mathbb{P}\widehat{\delta}_n^2) \}^{1/2} \lesssim M r_{m,n} \lesssim r_{m,n} \end{aligned}$$

both hold. Combining these results, we conclude that $|\hat{\delta}_t| \leq r_{m,n}$ on an event with probability at least $1 - Ce^{-p_n}$.

Next, we prove the second half of the proposition on the minimization condition (16). Denote the functions $\hat{\delta}_{\alpha,n} = \hat{\alpha}_n - \alpha_n^{\text{null}}, \, \delta_{\alpha,n,t} = \alpha_n^{\text{null}} - \alpha_{t,n}^*$ and $\hat{\delta}_{\alpha,n,t} = \hat{\alpha}_n - \alpha_{t,n}^* = \hat{\delta}_{\alpha,n} + \delta_{\alpha,n,t}$. It follows from the definition of $\hat{\delta}_t$ that $\mathbb{P}_n \ell(\cdot; \check{g}_n) - \mathbb{P}_n \ell(\cdot; \check{g}_n \pm r_{\inf,n} \hat{\alpha}_n) \leq 0$. We then have, for all $t \in \mathbb{R}$,

$$\begin{aligned} \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}) &- \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\alpha_{t,n}^{*}) \\ &= \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}) - \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\widehat{\alpha}_{n}) + \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\widehat{\alpha}_{n}) - \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\alpha_{t,n}^{*}) \\ &\leq \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\widehat{\alpha}_{n}) - \mathbb{P}_{n}\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\alpha_{t,n}^{*}) \\ &= \mathbb{P}_{n}\ell(\cdot;\widehat{g}_{n}+\widehat{\delta}_{t}\widehat{\alpha}_{n}\pm r_{\mathrm{inf},n}\widehat{\alpha}_{n}) - \mathbb{P}_{n}\ell(\cdot;\widehat{g}_{n}+\widehat{\delta}_{t}\widehat{\alpha}_{n}\pm r_{\mathrm{inf},n}\alpha_{t,n}^{*}) \\ &= \mathbb{P}_{n}\ell(\cdot;\widehat{g}_{n}+(\widehat{\delta}_{t}\pm r_{\mathrm{inf},n})\widehat{\alpha}_{n}) - \mathbb{P}_{n}\ell(\cdot;\widehat{g}_{n}+(\widehat{\delta}_{t}\pm r_{\mathrm{inf},n})\widehat{\alpha}_{n}\mp r_{\mathrm{inf},n}\widehat{\delta}_{\alpha,n,t}) \equiv J_{n,t}, \end{aligned}$$

where the inequality follows by our earlier remark.

Next,

$$\begin{split} J_{n,t} &= \frac{1}{n} \sum_{i=1}^{n} [\{\epsilon_{i} - \hat{\delta}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) \hat{\alpha}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j})\}^{2} \\ &- \{\epsilon_{i} - \hat{\delta}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) \hat{\alpha}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \pm r_{\mathrm{inf},n} \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j})\}^{2}] \\ &= \frac{1}{n} \sum_{i=1}^{n} \left[\mp 2 \left\{ \epsilon_{i} - \hat{\delta}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) \hat{\alpha}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \right\} r_{\mathrm{inf},n} \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - r_{\mathrm{inf},n}^{2} \hat{\delta}_{\alpha,n,t}^{2}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \right] \\ &= \pm (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) \frac{2r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \hat{\alpha}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \mp \frac{2r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \epsilon_{i} \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \\ &\pm \frac{2r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \hat{\delta}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - \frac{r_{\mathrm{inf},n}^{2}}{n} \sum_{i=1}^{n} \hat{\delta}_{\alpha,n,t}^{2}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \\ &\pm (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) \frac{2r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \alpha_{t,n}^{*}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) - \frac{r_{\mathrm{inf},n}^{2}}{n} \sum_{i=1}^{n} \hat{\delta}_{\alpha,n,t}^{2}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \\ &\pm \frac{2r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \hat{\delta}_{n}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \hat{\delta}_{\alpha,n,t}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) + (\pm 2\hat{\delta}_{t} + r_{\mathrm{inf},n}) \frac{r_{\mathrm{inf},n}}{n} \sum_{i=1}^{n} \hat{\delta}_{\alpha,n,t}^{2}(\widetilde{\mathbf{F}}_{i}, X_{i,j}) \\ &\pm 2r_{\mathrm{inf},n} (\hat{\delta}_{t} \pm r_{\mathrm{inf},n}) J_{n,t,1} \mp 2r_{\mathrm{inf},n} J_{n,t,2} + 2r_{\mathrm{inf},n} J_{n,t,3} + r_{\mathrm{inf},n} (\pm 2\hat{\delta}_{t} + r_{\mathrm{inf},n}) J_{n,t,4} \end{split}$$

where we have introduced

$$J_{n,t,1} = \frac{1}{n} \sum_{i=1}^{n} \alpha_{t,n}^{*}(\widetilde{F}_{i}, X_{i,j}) \widehat{\delta}_{\alpha,n,t}(\widetilde{F}_{i}, X_{i,j}), \quad J_{n,t,2} = \frac{1}{n} \sum_{i=1}^{n} \epsilon_{i} \widehat{\delta}_{\alpha,n,t}(\widetilde{F}_{i}, X_{i,j}), \\ J_{n,t,3} = \frac{1}{n} \sum_{i=1}^{n} \widehat{\delta}_{n}(\widetilde{F}_{i}, X_{i,j}) \widehat{\delta}_{\alpha,n,t}(\widetilde{F}_{i}, X_{i,j}), \quad J_{n,t,4} = \frac{1}{n} \sum_{i=1}^{n} \widehat{\delta}_{\alpha,n,t}^{2}(\widetilde{F}_{i}, X_{i,j}).$$

We only demonstrate the treatment for $J_{n,t,1}$ in detail; the treatments for the other $J_{n,k}$'s are in principle similar and have been employed in the proofs of Theorems 3.1, 3.2 and Proposition 3.3. For the term $J_{n,t,1}$, we further define $J_{n,t,1,1} = \mathbb{P}_n \alpha_{t,n}^* \hat{\delta}_{\alpha,n}$ and $J_{n,t,1,2} = \mathbb{P}_n \alpha_{t,n}^* \delta_{\alpha,n,t}$ so $J_{n,t,1} = J_{n,t,1,1} + J_{n,t,1,2}$. Now, we apply (B.15) in Proposition B.7 to conclude that with probability at least $1 - Ce^{-p_n}$,

$$\forall t \in \mathcal{T}_{\delta}, |J_{n,t,1,1}| = |(\mathbb{P}_n - \mathbb{P})\alpha_{t,n}^* \widehat{\delta}_{\alpha,n} + \mathbb{P}\alpha_{t,n}^* \widehat{\delta}_{\alpha,n}| \lesssim \{ \|\alpha_{t,n}^* \widehat{\delta}_{\alpha,n}\|_{L_2} + \nu_n \} \nu_n + \{ \mathbb{P}\alpha_{t,n}^{*2} \mathbb{P}\widehat{\delta}_{\alpha,n}^2 \}^{1/2}$$
$$\lesssim \|\widehat{\delta}_{\alpha,n}\|_{L_2} \lesssim r_{\alpha,n}^{\text{null}}.$$

To treat $J_{n,t,1,2}$, we apply the same proposition to conclude that with probability at least $1 - Ce^{-p_n}$,

$$\forall t \in \mathcal{T}_{\delta}, \quad |J_{n,t,1,2}| = |(\mathbb{P}_n - \mathbb{P})\alpha_{t,n}^*\delta_{\alpha,n,t} + \mathbb{P}\alpha_{t,n}^*\delta_{\alpha,n,t}| \lesssim \left\{ \|\alpha_{t,n}^*\delta_{\alpha,n,t}\|_{L_2} + \nu_n \right\} \nu_n + \left\{ \mathbb{P}\alpha_{t,n}^{*2}\mathbb{P}\delta_{\alpha,n,t}^2 \right\}^{1/2} \\ \lesssim \nu_n + r_{\alpha,t,b}.$$

Thus, with probability at least $1 - Ce^{-p_n}$,

$$\forall t \in \mathcal{T}_{\delta}, \quad |J_{n,t,1}| \lesssim r_{\alpha,n}^{\text{null}} + r_{\alpha,t,b}.$$

For the other $J_{n,t,k}$ terms, we can similarly conclude that, with probability at least $1 - Ce^{-p_n}$,

$$\forall t \in \mathcal{T}_{\delta}, \quad |J_{n,t,2}| \lesssim r_{m,n}(r_{\alpha,n}^{\text{null}} + r_{\alpha,t,b}), \quad |J_{n,t,3}| \lesssim r_{m,n}(r_{\alpha,n}^{\text{null}} + r_{\alpha,t,b}), \quad |J_{n,t,4}| \lesssim r_{\alpha,n}^{\text{null}}(r_{\alpha,n}^{\text{null}}\nu_n + r_{\alpha,t,b}).$$

Combining all the $J_{n,t,k}$ terms and further simplifying, we conclude that with probability at least $1 - Ce^{-p_n}$,

$$\forall t \in \mathcal{T}_{\delta}, \quad |J_{n,t}| \lesssim r_{\inf,n} r_{m,n} (r_{\alpha,n}^{\text{null}} + r_{\alpha,t,b}).$$

This concludes the proof of the proposition.

C.4 Proof of Theorem 3.4

Proof. As in the proof of Theorem 3.1, given (30), the proof of (31) is straightforward, so we will only prove (30). By Proposition 3.3 and Assumption 3.6, on the event \mathcal{A}_t (see Proposition 3.3) we have $|\hat{\delta}_t| \leq 1$, and hence further by (14), on the same event,

$$\check{g}_n \in \overline{\mathcal{F}}_n(\overline{r}+1) \equiv \{g + \delta_{\mathrm{t}}\alpha : g, \alpha \in \mathcal{F}_n(\overline{r}+1), |\delta_{\mathrm{t}}| \leq 1\}.$$

We now control the complexity of the class $\mathcal{F}_n(\bar{r}+1)$. Let Q be a generic measure on $(Y, \mathbf{F}, \mathbf{X})$, let \mathcal{N} be a cover of $\mathcal{F}_n(\bar{r}+1)$ by balls of radius $\tau/\sqrt{9}$ in the $L_2(Q)$ norm and moreover \mathcal{N} consist of functions uniformly bounded in magnitude by M, and let \mathcal{E} be a cover of the interval [-1, 1] by intervals of length at most $\tau/(\sqrt{9}M)$; we choose \mathcal{E} to have cardinality at most $[2\sqrt{9}M/\tau]$.

Suppose g, g', α, α' be functions of $(Y, \mathbf{F}, \mathbf{X})$, that are uniformly bounded in magnitude by M, and moreover satisfy $||g - g'||_{L_2(Q)} \leq \tau/\sqrt{9}$ and $||\alpha - \alpha'||_{L_2(Q)} \leq \tau/\sqrt{9}$, and let $\delta_t, \delta'_t \in [-1, 1]$ be such that $|\delta_t - \delta'_t| \leq \tau/(\sqrt{9}M)$. Let $\theta = g + \delta_t \alpha$ and $\theta' = g' + \delta'_t \alpha'$. Then

$$\begin{aligned} \|\theta' - \theta\|_{L_2(Q)}^2 &\leqslant 3\|g - g'\|_{L_2(Q)}^2 + 3\delta_t^2\|\alpha - \alpha'\|_{L_2(Q)}^2 + 3(\delta_t - \delta_t')^2\|\alpha'\|_{L_2(Q)}^2 \\ &\leqslant 3\|g - g'\|_{L_2(Q)}^2 + 3\delta_t^2\|\alpha - \alpha'\|_{L_2(Q)}^2 + 3(\delta_t - \delta_t')^2M^2 \leqslant 3\frac{\tau^2}{9} + 3\frac{\tau^2}{9} + 3\frac{\tau^2}{9} = \tau^2. \end{aligned}$$

Thus $N(\overline{\mathcal{F}}_n(\overline{r}+1), L_2(Q), \tau) \leq N(\mathcal{F}_n(\overline{r}+1), L_2(Q), \tau/\sqrt{9})^2 [2\sqrt{9}M/\tau]$ where we follow the notation for the covering numbers, for instance $N(\overline{\mathcal{F}}_n(\overline{r}+1), L_2(Q), \tau)$, in (A.1) in Proposition B.6. Thus, the logarithm of the covering number $N(\overline{\mathcal{F}}_n(\overline{r}+1), L_2(Q), \tau)$ for the class $\overline{\mathcal{F}}_n(\overline{r}+1)$ is bounded up to a multiplicative constant by the logarithm of the covering number for the class $\mathcal{F}_n(\overline{r}+1)$. Then, the proof of (30) in Theorem 3.4 follows that of (22) in Theorem 3.1 with minor modifications.

C.5 Proof of the size under the null: Theorem 3.5

Let $r = r(\cdot; b)$ be an unspecified, possibly random function that may change for each occurrence, but that always satisfies $\sup_{t \in \mathcal{T}_{\delta}} |r(t; b)| \leq b$. By this convention, the function $r = r(t; \mathcal{O}(1))$ in (C.19) satisfies $\sup_{t \in \mathcal{T}_{\delta}} |r(t; \mathcal{O}(1))| = \mathcal{O}(1)$.

Proof of Theorem 3.5. We will first prove the following master approximate i.i.d. expansion for the statistic $\check{\eta}_t^{s}(\check{g}_n)$: under the null hypothesis H_0 in (3), with probability at least $1 - Ce^{-p_n}$, uniformly over all $t \in \mathcal{T}_{\delta}$ we have

$$\sqrt{n}\widetilde{\eta}_{t}^{s}(\widetilde{g}_{n}) = \frac{t}{\sqrt{n}}\sum_{i=1}^{n}\widetilde{\epsilon}_{i}\alpha_{n}^{\text{null}}(\widetilde{F}_{i}, X_{i,j}) + r(t; \mathcal{O}(1)).$$
(C.19)

Recall the notations $\check{\eta}_t^{s}$ from (17), $\tilde{\eta}_t^{s}$ from (24), and the random variable $Z_{t,j,h}$ from (26). Introduce, for a generic function $m : \mathbb{R}^{r+1} \to \mathbb{R}$, the functional η_t^{s} as

$$\eta_t^{\mathrm{s}}(m) = \int_{\Omega_h} \left[\exp\{t \, m_j^{\mathrm{s}}(\boldsymbol{f}, x_j)\} - 1 \right] \mathrm{d}\mathbb{P}.$$

Under either H_0 and H_1 , we have the master decomposition

$$\widetilde{\eta}_{t}^{s}(\widetilde{g}_{n}) = \{\widetilde{\eta}_{t}^{s}(\widetilde{g}_{n}) - \widetilde{\eta}_{t}^{s}(\widetilde{g}_{n})\} + \{\widetilde{\eta}_{t}^{s}(\widetilde{g}_{n}) - \widetilde{\eta}_{t}^{s}(g_{0})\} + \{\widetilde{\eta}_{t}^{s}(g_{0}) - \eta_{t}^{s}(m_{0})\} \\
+ \{\eta_{t}^{s}(m_{0}) - \int_{\Omega_{h}} \left[\exp\{t \, m_{0,j}(\boldsymbol{f}, x_{j})\} - 1\right] d\mathbb{P}\} + \int_{\Omega_{h}} \left[\exp\{t \, m_{0,j}(\boldsymbol{f}, x_{j})\} - 1\right] d\mathbb{P} \\
\equiv I_{t,n,1} + I_{t,n,2} + \underbrace{I_{t,n,3} + I_{t,n,4} + I_{t,n,5}}_{=\mathbb{E}Z_{t,j,h}} .$$
(C.20)

In this decomposition,

- i) $I_{t,n,1}$ is a centered empirical process evaluated at the random function \check{g}_n . Under the null, $I_{t,n,1}$ admits a fast rate converging to zero because here the smoothed derivative estimator $\check{g}_{n,j}^{s}$ is expected to be close to zero. Under the alternative, $I_{t,n,1}$ still converges to zero due to the averaging effect, though the rate will be comparatively slower.
- ii) $I_{t,n,2}$ is the dominant term in the decomposition under the null, and in fact will admit an approximate i.i.d. expansion that will in turn become the leading, i.i.d. sum in (C.19).
- iii) The term $I_{t,n,3}$ represents the approximation error arising from replacing the latent factor \mathbf{F} by the diversified factor $\mathbf{\tilde{F}}$, $I_{t,n,4}$ is a bias term arising from smoothing the true regression function derivative $m_{0,j}$, and finally, $I_{t,n,5}$ can be regarded as the original signal strength favoring the alternative. Together, the (magnitude of the) sum $I_{t,n,3}+I_{t,n,4}+I_{t,n,5}$ can be regarded as containing the signal favoring the alternative. Here, under the null, the sum becomes $\tilde{\eta}_t^s(g_0) = \mathbb{E}Z_{t,j,h}$ which is zero (see the remark below (26)).

Next, the treatments of the terms $I_{t,n,1}$ and $I_{t,n,2}$ will be covered in Sections C.5.1 and C.5.2 respectively. The transition from (32) to (33) will be covered in Section C.5.3 by showing that the estimated variance $\|\hat{\epsilon}\hat{\alpha}_n\|_{L_2(\mathbb{P}_n)}^2$ is a consistent estimator of the unknown asymptotic variance $\|\epsilon\alpha_n^{\text{null}}\|_{L_2}^2$. We finalize our proof of Theorem 3.5 afterward.

C.5.1 Proof for the first term in (C.20)

Now we handle the term $I_{t,n,1}$ in the decomposition (C.20), which as stated below (C.20) is a centered empirical process term. Recall the definition of $\mathcal{T}_{\delta} = [-T, -\delta] \cup [\delta, T]$.

Proposition C.1. Suppose that the conditions of Theorem 3.5 hold. Then, under either H_0 or H_1 , with probability at least $1 - Ce^{-p_n}$, for a constant $c_{4,1}$, and uniformly over $t \in \mathcal{T}_{\delta}$,

$$I_{t,n,1} \leqslant c_{4,1} \{ \frac{|t|}{h} r_{m,n} + \frac{1}{h} \nu_n + ||Z_{t,j,h}||_{L_2} \} \nu_n.$$

Proof. Define the function $\theta(\cdot; g, t) = \theta(\tilde{f}, x_j; g, t) = \{\exp[t \Psi\{g_j^{s}(\tilde{f}, x_j)\}] - 1\}\mathbb{1}\{x_j \in \mathcal{B}_h\} : \mathbb{R}^{\bar{r}+1} \to \mathbb{R} \text{ (essentially the summand in (17)) indexed by } g \in \overline{\mathcal{F}}_n(\bar{r}+1) \text{ and } t.$ Then, $I_{t,n,1} = (\mathbb{P}_n - \mathbb{P})\theta(\cdot; \check{g}_n, t)$. Now, in the proof of Theorem 3.4 in Section C.4, it was shown

that with probability at least $1 - Ce^{-p_n}$, \check{g}_n belongs to the functional class $\overline{\mathcal{F}}_n(\bar{r}+1)$ defined in that proof. We shall focus on this event.

As in the proof of Theorem 3.4, let g, g' be functions of $(Y, \mathbf{F}, \mathbf{X})$, and moreover let $t, t' \in \mathcal{T}_{\delta}$. Then

$$\begin{aligned} |\theta(\widetilde{\boldsymbol{f}}, x_j; g, t) - \theta(\widetilde{\boldsymbol{f}}, x_j; g', t')| &= |\exp[t \,\Psi\{g_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j)\}] - \exp[t' \,\Psi\{g_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j)\}] |\mathbbm{1}\{x_j \in \mathcal{B}_h\} \\ &\lesssim t |(g - g')_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j)| \mathbbm{1}\{x_j \in \mathcal{B}_h\} + |t - t'| \leqslant T |(g - g')_j^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j)| \mathbbm{1}\{x_j \in \mathcal{B}_h\} + |t - t'|. \end{aligned}$$

Since $t \in \mathcal{T}_{\delta}$ and \mathcal{T}_{δ} is compact, the class of functions $\{\theta(\cdot; g, t) : g \in \overline{\mathcal{F}}_n(\overline{r} + 1), t \in \mathcal{T}_{\delta}\}$ admits a uniform covering number similar to that in (B.22). Thus, the situation is largely similar to that in Proposition B.8, and we obtain that on an event with probability at least $1 - 12e^{-s}$, for a constant c'_6 ,

$$\forall g \in \overline{\mathcal{F}}_n(\overline{r}+1), \forall t \in \mathcal{T}_{\delta}, \quad |(\mathbb{P}_n - \mathbb{P})\theta(\cdot; g, t)| \leq c_6' [\{\|\theta\|_{L_2} + \nu_{\mathrm{s},n}\} \{\nu_{\mathrm{s},n} + \sqrt{\frac{s}{n}}\} + \frac{1}{h^2}\nu_{\mathrm{s},n}^2].$$

Choosing s as in the proof in Section C.1 of Theorem 3.1 in the above, we obtain that, with probability at least $1 - Ce^{-p_n}$, by changing the constant c'_6 if necessary,

$$\forall t \in \mathcal{T}_{\delta}, \quad I_{t,n,1} \leqslant c_{6}' \{ \|\theta(\cdot; \check{g}_{n}, t)\|_{L_{2}} + \frac{1}{h}\nu_{n} \} \nu_{n}.$$
(C.21)

Now, with the aid of (A.2) and the definition of $Z_{t,j,h}$ from (26),

$$\begin{split} \|\theta(\cdot;\check{g}_{n},t)\|_{L_{2}}^{2} &= \int_{\Omega_{h}} \{\exp(t\,\Psi\{\check{g}_{n,j}^{s}(\tilde{f},x_{j})\}) - 1\}^{2} \mathrm{d}\mathbb{P} \\ &\leq 2\int_{\Omega_{h}} \{\exp(t\,\Psi\{\check{g}_{n,j}^{s}(\tilde{f},x_{j})\}) - \exp(t\,\Psi\{g_{0,j}^{s}(\tilde{f},x_{j})\})\}^{2} \mathrm{d}\mathbb{P} + 2\int_{\Omega_{h}} \{\exp(t\,\Psi\{g_{0,j}^{s}(\tilde{f},\mathrm{d}x_{j})\}) - 1\}^{2} \mathrm{d}\mathbb{P} \\ &\lesssim t^{2}\int_{\Omega_{h}} \{(\check{g}_{n}-g_{0})_{j}^{s}\}^{2} \mathrm{d}\mathbb{P} + \|Z_{t,j,h}\|_{L_{2}}^{2} \lesssim \frac{t^{2}}{h^{2}}\int(\check{g}_{n}-g_{0})^{2} \mathrm{d}\mathbb{P} + \|Z_{t,j,h}\|_{L_{2}}^{2} \\ &\lesssim \frac{t^{2}}{h^{2}}\int\{(\check{g}_{n}-m_{0})^{2} + (g_{0}-m_{0})^{2}\} \mathrm{d}\mathbb{P} + \|Z_{t,j,h}\|_{L_{2}}^{2} \lesssim \frac{t^{2}}{h^{2}}r_{m,n}^{2} + \|Z_{t,j,h}\|_{L_{2}}^{2}, \end{split}$$

where the last step follows by Theorem 3.4 and Lemma B.2. Combining the above with (C.21) then yield the conclusion of the proposition.

C.5.2 Proof for the second term in (C.20)

Recall that $\alpha_{t,n}^* = \alpha_{t,n}^*(\widetilde{f}, x_j) : \mathbb{R}^{\overline{r}+1} \to \mathbb{R}$ is the Riesz representer that satisfies (25).

Proposition C.2. Suppose that the conditions of Theorem 3.5 hold. Then, under either H_0 or H_1 , with probability at least $1 - Ce^{-p_n}$, for a common function $r = r(\cdot; b)$ as introduced above Theorem 3.5, uniformly over $t \in \mathcal{T}_{\delta}$ we have

$$I_{t,n,2} = t_{\overline{n}}^{1} \sum_{i=1}^{n} \widetilde{\epsilon}_{i} \alpha_{t,n}^{*} (\widetilde{F}_{i}, X_{i,j}) + |t| r(t; r_{m,n} r_{\alpha,t,n} + \delta_{f}) + (|t| + t^{2}) r(t; h^{-2} r_{m,n}^{2}).$$
(C.22)

Proof. To streamline our presentation, we divide the proof into two parts. In the first part, Lemma C.3, the proposition is proven under a set of technical conditions, while in the second part, Lemma C.4, the technical conditions themselves are verified. The general direction of the proof of Lemma C.3 follows that of Chen et al. (2014, 2022), but here we work with the tweaked regression function estimator \check{g}_n instead of the preliminary estimator \hat{g}_n , and also take into account the accuracy loss of using \tilde{F} in place of F. More importantly, by using the tweaked estimator \check{g}_n , we avoid artificially imposing the

"undersmoothing/bias vanishing abnormally fast" condition as in Chen et al. (2022); more specifically, see how their "undersmoothing" condition, Assumption 4.4, is invoked in the proof of their Lemma C.1(ii). If we were to impose a similar condition, then we would need $\nu_{a,n} = \nu_{a,n,\text{undersmooth}}$ from (C.11), the bound on the bias, to be $\mathcal{O}(n^{-1/2})$; this would lead to a scaling of $\widetilde{L}\widetilde{k}_0$ that would further result in $\nu_{s,n} = \nu_{s,n,\text{undersmooth}}$ from (B.11), the accompanying standard deviation for the stochastic variation, to be at least on the order of $n^{-\frac{2\kappa-1}{4\kappa}}$ (up to log factors). Clearly, the rates for $\nu_{a,n,\text{undersmooth}}$ and $\nu_{s,n,\text{undersmooth}}$ in the bias/variance decomposition are no longer balanced, and in particular the rate for $\nu_{s,n,\text{undersmooth}}$ is inferior to our balanced rate ν_n in (21); in fact, when $\kappa < 1/2$, it is no longer true that $\nu_{s,n,\text{undersmooth}} = \mathcal{O}(1)$, so the regression function estimators \widehat{g}_n and \widecheck{g}_n are no longer consistent.

Lemma C.3. Suppose that Assumption 3.6(ii) holds, and that on an event \mathcal{A} , for a common function $r = r(\cdot; b)$ as introduced above Theorem 3.5, uniformly over $t \in \mathcal{T}_{\delta}$ the following conditions hold:

$$\mathbb{P}_n\ell(\cdot;\check{g}_n) - \mathbb{P}_n\ell(\cdot;\check{g}_n \pm r_{\inf,n}\alpha^*_{t,n}) \leqslant r(t;r_{\inf,n}r_{m,n}r_{\alpha,t,n})$$
(C.23)

(note that (C.23) is essentially the minimization condition (16) at t with $g = \check{g}_n$ for the tolerance on the order of $r_{m,n}(r_{\alpha,n}^{\text{null}} + r_{\alpha,t,b})$),

$$\mathbb{P}\{\ell(\cdot;\check{g}_n) - \ell(\cdot;\check{g}_n \pm r_{\inf,n}\alpha_{t,n}^*)\} \\ \pm 2r_{\inf,n}\langle\check{g}_n - m_0, \alpha_{t,n}^*\rangle \mp 2r_{\inf,n}\mathbb{E}[\epsilon\alpha_{t,n}^*(\widetilde{F}, X_j)] = r(t; r_{\inf,n}^2),$$
(C.24)

$$\left(\mathbb{P}_{n}-\mathbb{P}\right)\left\{\ell(\cdot;\check{g}_{n})-\ell(\cdot;\check{g}_{n}\pm r_{\mathrm{inf},n}\alpha_{t,n}^{*})\mp 2r_{\mathrm{inf},n}\epsilon\alpha_{t,n}^{*}\right\}=r(t;r_{\mathrm{inf},n}r_{m,n}\nu_{n}),\qquad(C.25)$$

$$\widetilde{\eta}_t^{\mathbf{s}}(\widetilde{g}_n) - \widetilde{\eta}_t^{\mathbf{s}}(g_0) - \frac{\partial \widetilde{\eta}_t^{\mathbf{s}}(g_0)}{\partial g} [\widetilde{g}_n - g_0] = (|t| + t^2) r(t; h^{-2} r_{m,n}^2), \quad (C.26)$$

$$\frac{\partial \tilde{\eta}_t^{\mathrm{s}}(g_0)}{\partial g} [\tilde{g}_n - g_0] = t \langle \tilde{g}_n - g_0, \alpha_{t,n}^* \rangle.$$
(C.27)

Then, on the intersection of the event \mathcal{A} and another event with probability at least $1 - Ce^{-p_n}$, uniformly over $t \in \mathcal{T}_{\delta}$, (C.22) in Proposition C.2 hold.

Proof. By condition (C.25),

$$\mathbb{P}_{n}\ell(\cdot;\check{g}_{n}) - \mathbb{P}_{n}\ell(\cdot;\check{g}_{n} \pm r_{\inf,n}\alpha_{t,n}^{*}) = \mathbb{P}\{\ell(\cdot;\check{g}_{n}) - \ell(\cdot;\check{g}_{n} \pm r_{\inf,n}\alpha_{t,n}^{*})\} \pm 2r_{\inf,n}(\mathbb{P}_{n} - \mathbb{P})\{\epsilon\alpha_{t,n}^{*}\} \\
+ (\mathbb{P}_{n} - \mathbb{P})\{\ell(\cdot;\check{g}_{n}) - \ell(\cdot;\check{g}_{n} \pm r_{\inf,n}\alpha_{t,n}^{*}) \mp 2r_{\inf,n}\epsilon\alpha_{t,n}^{*}\} \\
= \mathbb{P}\{\ell(\cdot;\check{g}_{n}) - \ell(\cdot;\check{g}_{n} \pm r_{\inf,n}\alpha_{t,n}^{*})\} \pm 2r_{\inf,n}(\mathbb{P}_{n} - \mathbb{P})\{\epsilon\alpha_{t,n}^{*}\} + r(t;r_{\inf,n}r_{m,n}\nu_{n}).$$

Combining the above with condition (C.24) and Assumption 3.6(ii) yields

$$\mathbb{P}_n\ell(\cdot;\check{g}_n) - \mathbb{P}_n\ell(\cdot;\check{g}_n \pm r_{\inf,n}\alpha_{t,n}^*) = \mp 2r_{\inf,n}\langle\check{g}_n - m_0, \alpha_{t,n}^*\rangle \pm 2r_{\inf,n}\mathbb{P}_n\{\epsilon\alpha_{t,n}^*\} + r(t;r_{\inf,n}r_{m,n}\nu_n)$$

By the above and the minimization condition (C.23), we further obtain

$$\langle \check{g}_n - m_0, \alpha_{t,n}^* \rangle - \mathbb{P}_n \{ \epsilon \alpha_{t,n}^* \} = r(t; r_{m,n} r_{\alpha,t,n}).$$
(C.28)

Next, we show that $t\langle \check{g}_n - m_0, \alpha_{t,n}^* \rangle$ can approximate $\eta_t^{s}(\check{g}_n) - \eta_t^{s}(g_0)$ well. First, by conditions (C.26) and (C.27),

$$I_{t,n,2} = \widetilde{\eta}_t^{\mathrm{s}}(\widecheck{g}_n) - \widetilde{\eta}_t^{\mathrm{s}}(g_0) = \frac{\partial \widetilde{\eta}_t^{\mathrm{s}}(g_0)}{\partial g} [\widecheck{g}_n - g_0] + \left\{ \widetilde{\eta}_t^{\mathrm{s}}(\widecheck{g}_n) - \widetilde{\eta}_t^{\mathrm{s}}(g_0) - \frac{\partial \widetilde{\eta}_t^{\mathrm{s}}(g_0)}{\partial g} [\widecheck{g}_n - g_0] \right\}$$
$$= t \langle \widecheck{g}_n - g_0, \alpha_{t,n}^* \rangle + (|t| + t^2) r(t; h^{-2} r_{m,n}^2).$$

Combining the above with (C.28), and using the Cauchy-Schwarz inequality followed by Lemma B.2, we obtain

$$\begin{split} I_{t,n,2} &= t \langle \widetilde{g}_n - g_0, \alpha_{t,n}^* \rangle + (|t| + t^2) r(t; h^{-2} r_{m,n}^2) \\ &= t \, r(t; r_{m,n} r_{\alpha,t,n}) + (|t| + t^2) r(t; h^{-2} r_{m,n}^2) + t \langle m_0 - g_0, \alpha_{t,n}^* \rangle + t \frac{1}{n} \sum_{i=1}^n \epsilon_i \alpha_{t,n}^* (\widetilde{F}_i, X_{i,j}) \\ &= t \, r(t; r_{m,n} r_{\alpha,t,n} + \delta_{\rm f}) + (|t| + t^2) r(t; h^{-2} r_{m,n}^2) + t \frac{1}{n} \sum_{i=1}^n \epsilon_i \alpha_{t,n}^* (\widetilde{F}_i, X_{i,j}). \end{split}$$

Compared to (C.22), we only need to replace the ϵ_i 's by the $\tilde{\epsilon}_i$'s in the last line above. This can be done similarly to how we dealt with the term $I_{n,2,2}$ in the proof of Theorem 3.1 in Section C.1, and is why the intersection of \mathcal{A} with another event with probability at least $1 - Ce^{-p_n}$ in the statement of Lemma C.3 is needed, at the cost of another $tr(t; \delta_f)$ term. Eventually we arrive at

$$I_{t,n,2} - t\frac{1}{n}\sum_{i=1}^{n} \widetilde{\epsilon}_{i} \alpha_{t,n}^{*}(\widetilde{F}_{i}, X_{i,j}) = t r(t; r_{m,n}r_{\alpha,t,n} + \delta_{f}) + (|t| + t^{2})r(t; h^{-2}r_{m,n}^{2}),$$

which is the conclusion of the lemma.

Lemma C.4. Suppose that the conditions of Theorem 3.5 hold. Then, conditions (C.23), (C.24), (C.25), (C.26) and (C.27) hold with probability at least $1 - Ce^{-p_n}$.

Proof. As stated in Lemma C.3, Ineq. (C.23) is essentially the minimization condition (16), which has already been established in Proposition 3.3. Next, we verify condition (C.24). We have

$$\ell(Y, \widetilde{\boldsymbol{F}}, X_j; \widecheck{g}_n) - \ell(Y, \widetilde{\boldsymbol{F}}, X_j; \widecheck{g}_n \pm r_{\text{inf},n} \alpha_{t,n}^*) = \{Y - \widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j)\}^2 - \{Y - \widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j) \mp r_{\text{inf},n} \alpha_{t,n}^*(\widetilde{\boldsymbol{F}}, X_j)\}^2
= [\epsilon - \{\widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j) - m_0(\boldsymbol{F}, X_j)\}]^2 - [\epsilon - \{\widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j) - m_0(\boldsymbol{F}, X_j)\} \mp r_{\text{inf},n} \alpha_{t,n}^*(\widetilde{\boldsymbol{F}}, X_j)]^2
= \pm 2r_{\text{inf},n}[\epsilon - \{\widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j) - m_0(\boldsymbol{F}, X_j)\}] \alpha_{t,n}^*(\widetilde{\boldsymbol{F}}, X_j) - r_{\text{inf},n}^2 \alpha_{t,n}^{*2}(\widetilde{\boldsymbol{F}}, X_j)
= \mp 2r_{\text{inf},n}\{\widecheck{g}_n(\widetilde{\boldsymbol{F}}, X_j) - m_0(\boldsymbol{F}, X_j)\} \alpha_{t,n}^*(\widetilde{\boldsymbol{F}}, X_j) - r_{\text{inf},n}^2 \alpha_{t,n}^{*2}(\widetilde{\boldsymbol{F}}, X_j) \pm 2r_{\text{inf},n} \epsilon \alpha_{t,n}^*(\widetilde{\boldsymbol{F}}, X_j).$$
(C.29)

Then, (C.24) follows by taking the integral of (C.29) against the measure \mathbb{P} .

Next, we verify condition (C.25). Starting from Eq. (C.29),

$$(\mathbb{P}_n - \mathbb{P}) \left\{ \ell(\cdot; \check{g}_n) - \ell(\cdot; \check{g}_n \pm r_{\inf,n} \alpha_{t,n}^*) \mp 2r_{\inf,n} \epsilon \alpha_{t,n}^* \right\}$$

= $\mp 2r_{\inf,n} (\mathbb{P}_n - \mathbb{P}) \{ (\check{g}_n - m_0) \alpha_{t,n}^* \} - r_{\inf,n}^2 (\mathbb{P}_n - \mathbb{P}) \alpha_{t,n}^{*2} = r(t; r_{\inf,n} r_{m,n} \nu_n),$

where the last step holds with probability at least $1 - Ce^{-p_n}$, and follows by (at this point) the usual application of (B.15) in Proposition B.7, with $r_{m,n}$ from Theorem 3.4.

Next, we verify condition (C.26). In condition (C.26), the second-order remainder term is, for some τ' strictly between 0 and 1, and $v = \check{g}_n - g_0$,

$$\begin{split} &\frac{1}{2} \left[\frac{\partial^2}{\partial \tau^2} \int_{\Omega_h} \left[e^{t \,\Psi\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\}} - 1 \right] \mathrm{d}\mathbb{P} \right] \Big|_{\tau = \tau'} \\ &= \frac{1}{2} \left[\frac{\partial}{\partial \tau} \left[\int_{\Omega_h} t \dot{\Psi}\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\} v_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j) e^{t \,\Psi\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\}} \mathrm{d}\mathbb{P} \right] \right] \Big|_{\tau = \tau'} \\ &= \frac{1}{2} \left[\int_{\Omega_h} \left[t \ddot{\Psi}\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\} (v_j^{\mathsf{s}})^2(\tilde{\boldsymbol{f}}, x_j) e^{t \,\Psi\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\}} + t^2 \dot{\Psi}^2\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\} (v_j^{\mathsf{s}})^2(\tilde{\boldsymbol{f}}, x_j) e^{t \,\Psi\{(g_0 + \tau v)_j^{\mathsf{s}}(\tilde{\boldsymbol{f}}, x_j)\}} \right] \mathrm{d}\mathbb{P} \right] \Big|_{\tau = \tau'}, \end{split}$$

which by Assumption B.5 and the boundedness of \mathcal{T}_{δ} is bounded in magnitude and up to a multiplicative factor (that is not dependent on t) by $(|t| + t^2) \int_{\Omega_h} v_j^{\mathrm{s}}(\tilde{f}, x_j)^2 d\mathbb{P}$. Thus, condition (C.26) is verified because, by the first half of (A.2) in Lemma A.1 applied to this result,

$$\begin{aligned} &|\widetilde{\eta}_{t}^{s}(\widetilde{g}_{n}) - \widetilde{\eta}_{t}^{s}(g_{0}) - \frac{\widetilde{\eta}_{t}^{s}(g_{0})}{\partial g} [\widetilde{g}_{n} - g_{0}]|_{L_{2}} \lesssim (|t| + t^{2}) ||(\widetilde{g}_{n} - g_{0})_{j}^{s}||_{L_{2}}^{2} \\ &\lesssim (|t| + t^{2}) \frac{1}{h^{2}} ||\widetilde{g}_{n} - g_{0}||_{L_{2}}^{2} \lesssim (|t| + t^{2}) \frac{1}{h^{2}} r_{m,n}^{2}. \end{aligned}$$

Finally, we verify condition (C.27). We just need to recall the first part of Lemma A.3, from which condition (C.27) follows as a special case. This completes the proof of Lemma C.4. \Box

Then, as mentioned earlier, combining Lemmata C.3 and C.4 completes the proof of Proposition C.2. $\hfill \Box$

C.5.3 Variance estimation

Proposition C.5. Suppose that the conditions of Theorem 3.5 hold. Then, under either H_0 or H_1 , with probability at least $1 - Ce^{-p_n}$,

$$\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}^2 - \|\epsilon\alpha_n^{\text{null}}\|_{L_2}^2 = \mathcal{O}(1)$$

Hence, under H_0 , we are free to replace $\|\epsilon \alpha_n^{\text{null}}\|_{L_2}^2$ in (32) by $\|\widehat{\epsilon} \widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}^2$ to arrive at (33).

Proof. We start from the overall decomposition

$$\begin{split} \|\widehat{\epsilon}\widehat{\alpha}_{n}\|_{L_{2}(\mathbb{P}_{n})}^{2} - \|\epsilon\alpha_{n}^{\text{null}}\|_{L_{2}}^{2} &= \frac{1}{n}\sum_{i=1}^{n}\{m_{0}(\boldsymbol{F}_{i}, X_{i,j}) + \epsilon_{i} - \check{g}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j})\}^{2}\widehat{\alpha}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j})^{2} - \|\epsilon\alpha_{n}^{\text{null}}\|_{L_{2}}^{2} \\ &= \underbrace{\frac{1}{n}\sum_{i=1}^{n}\{\check{g}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j}) - m_{0}(\boldsymbol{F}_{i}, X_{i,j})\}^{2}\widehat{\alpha}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j})^{2}}_{J_{n,1}} \\ &- 2\underbrace{\frac{1}{n}\sum_{i=1}^{n}\epsilon_{i}\{\check{g}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j}) - m_{0}(\boldsymbol{F}_{i}, X_{i,j})\}\widehat{\alpha}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j})^{2}}_{J_{n,2}} \\ &+ \underbrace{\frac{1}{n}\sum_{i=1}^{n}[\epsilon_{i}^{2}\widehat{\alpha}_{n}(\widetilde{\boldsymbol{F}}_{i}, X_{i,j})^{2} - \mathbb{P}\{\epsilon^{2}\widehat{\alpha}_{n}^{2}\}]}_{J_{n,3}} + \underbrace{\mathbb{P}\{\epsilon^{2}\widehat{\alpha}_{n}^{2}\} - \|\epsilon\alpha_{n}^{\text{null}}\|_{L_{2}}^{2}}_{J_{n,4}} \end{split}$$
$$= J_{n,1} - 2J_{n,2} + J_{n,3} + J_{n,4}$$

The term $J_{n,1}$ can be treated similarly to how we dealt with the term $I_{n,1}$ in the proof of Theorem 3.1 in Section C.1, specifically by invoking Lemma B.6 and the choice of s' in the proof of Theorem 3.1. We conclude that with probability at least $1 - Ce^{-p_n}$,

$$|J_{n,1}| \leq |(\mathbb{P}_n - \mathbb{P})[\{\check{g}_n - m_0\}^2 \widehat{\alpha}_n^2]| + \mathbb{P}[\{\check{g}_n - m_0\}^2 \widehat{\alpha}_n^2] \\ \lesssim \mathbb{P}[\{\check{g}_n - m_0\}^2 \widehat{\alpha}_n^2] + \nu_n^2 + \mathbb{P}[\{\check{g}_n - m_0\}^2 \widehat{\alpha}_n^2] \lesssim r_{m,n}^2 M^2 \lesssim r_{m,n}^2.$$

The term $J_{n,2}$ can be treated similarly to how we dealt with the term $I_{n,2}$ in the proof of Theorem 3.1. We conclude that with probability at least $1 - Ce^{-p_n}$,

$$|J_{n,2}| \lesssim \left\{ \| (\check{g}_n - m_0) \widehat{\alpha}_n^2 \|_{L_2} + \nu_n \right\} \nu_n + \delta_{\mathrm{f}} \| (\check{g}_n - m_0) \widehat{\alpha}_n^2 \|_{L_2} \lesssim r_{m,n}^2.$$

The term $J_{n,3}$ is a centered empirical process $(\mathbb{P}_n - \mathbb{P})\{\epsilon^2 \alpha^2\}$ evaluated at the random function $\alpha = \hat{\alpha}_n$. This time neither Lemma B.6 nor Proposition B.7 directly applies because the factor ϵ^2 involved is sub-exponential instead of sub-Gaussian. Nevertheless one can still apply Theorem 1 in Adamczak (2008) to conclude that with probability at least $1 - Ce^{-p_n}$, $|J_{n,3}| \leq \nu_n$.

Finally, for the term $J_{n,4}$, first using the sub-Gaussianity of ϵ and then the Cauchy-Schwarz inequality,

$$\begin{aligned} \|\mathbb{P}\{\epsilon^{2}\widehat{\alpha}_{n}^{2}\} - \|\epsilon\alpha_{n}^{\mathrm{null}}\|_{L_{2}}^{2} &\leq \|\mathbb{P}\{\widehat{\alpha}_{n}^{2} - (\alpha_{n}^{\mathrm{null}})^{2}\}\| = \|\mathbb{P}\{(\widehat{\alpha}_{n} + \alpha_{n}^{\mathrm{null}})(\widehat{\alpha}_{n} - \alpha_{n}^{\mathrm{null}})\}\| \\ &\lesssim [\mathbb{P}\{(\widehat{\alpha}_{n} + \alpha_{n}^{\mathrm{null}})^{2}\}\mathbb{P}\{(\widehat{\alpha}_{n} - \alpha_{n}^{\mathrm{null}})^{2}\}]^{1/2} \lesssim r_{\alpha,n}^{\mathrm{null}} \end{aligned}$$

where the last step holds with probability at least $1 - Ce^{-p_n}$ by Theorem 3.2.

Collecting terms, we conclude that on an event with probability at least $1 - Ce^{-p_n}$, $\|\|\widehat{\epsilon \alpha}_n\|_{L_2(\mathbb{P}_n)}^2 - \|\epsilon \alpha_n^{\text{null}}\|_{L_2}^2\| \lesssim r_{m,n}^2 + \nu_n + r_{\alpha,n}^{\text{null}}$ which is easily $\mathcal{O}(1)$ under condition (*) in Theorem 3.5. This completes the proof of the proposition. \Box

Now we are ready to put everything together to complete the proof of the component of Theorem 3.5 regarding the fixed-*t* test (18). Under the null, as argued in the proof of Theorem 3.5 in the main text, the sum $I_{t,n,3} + I_{t,n,4} + I_{t,n,5} = 0$ in (C.20). Next, (under the null) the leading term in (C.22) in Proposition C.2 becomes $t_n^{\frac{1}{n}} \sum_{i=1}^n \tilde{\epsilon}_i \alpha_n^{\text{null}} (\tilde{F}_i, X_{i,j})$, while the term $I_{t,n,1}$ and the remainder terms in (C.22) are bounded by $\mathcal{O}(n^{-1/2})$ with probability at least $1 - Ce^{-p_n}$ by Propositions C.1 and C.2 and the additional conditions in the statements of Theorem 3.5. This concludes the i.i.d. representation in (C.19). Next, (32) follows from (C.19) by an application of a central limit theorem followed by a straightforward replacement of $\|\tilde{\epsilon}\alpha_n^{\text{null}}\|_{L_2}$ by $\|\epsilon\alpha_n^{\text{null}}\|_{L_2}$ based on (B.1). Finally, the transition from (32) to (33) is covered by Proposition C.5 in Section C.5.3.

Finally, we prove the component of Theorem 3.5 regarding the square test involving $\hat{\chi}^2$ in (20); the proof for the sup test involving \hat{Z} in (19) is analogous. By our earlier results for the fixed-t test, (under the null and the conditions of that theorem) with probability at least $1 - Ce^{-p_n}$, uniformly over all $t \in \mathcal{T}_{\delta}$ we have

$$\frac{n}{t^2} \{ \breve{\eta}_t^{\mathrm{s}}(\breve{g}_n) \}^2 = \{ \frac{1}{\sqrt{n}} \sum_{i=1}^n \widetilde{\epsilon}_i \alpha_n^{\mathrm{null}}(\widetilde{F}_i, X_{i,j}) \}^2 + \{ \frac{1}{\sqrt{n}} \sum_{i=1}^n \widetilde{\epsilon}_i \alpha_n^{\mathrm{null}}(\widetilde{F}_i, X_{i,j}) \} r(t; \mathcal{O}(1)) + r(t; \mathcal{O}(1))^2 \}$$

Therefore, by the central limit theorem and the Slutsky's theorem,

$$\int_{\mathcal{T}_{\delta}} \frac{n}{t^2} \{ \check{\eta}_t^{\mathrm{s}}(\check{g}_n) \}^2 \to_d \left\{ \frac{1}{T} \int_{\mathcal{T}_{\delta}} w(t) \mathrm{d}t \right\} \| \check{\epsilon} \alpha_n^{\mathrm{null}} \|_{L_2} \chi_1^2$$

which is equivalent to the last equation in (32). Again by Proposition C.5 and the Slutsky's theorem, we are further free to replace $\|\widetilde{\epsilon}\alpha_n^{\text{null}}\|_{L_2}$ in (32) by $\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}$.

C.6 Power of the fixed-t test under the alternative

In this subsection, our principal Theorem C.6 guarantees that the power of our fixed-t test in (18) approaches one under all local alternatives satisfying conditions (i) and (ii) in the theorem that place mild conditions on how fast $|\mathbb{E}Z_{t,j,h}|$, the signal strength favoring the alternative (as explained in the proof of Theorem 3.5), can decay. Similar results for the ensemble tests (19) and (20) can be shown using analogous techniques and we omit the details here.

Theorem C.6. Suppose that the conditions of Theorem 3.5 hold. Furthermore, suppose that at some $t \in \mathcal{T}_{\delta}$, for large enough constants $C_{\text{alt},1}$ and $C_{\text{alt},2}$, (i) $C_{\text{alt},1}\sqrt{\log(n)/n} \leq |\mathbb{E}Z_{t,j,h}|$; (ii) $C_{\text{alt},2}\{\delta_{\text{f}} + \frac{1}{h^2}(\nu_n^2 + \delta_{\text{f}}^2) + \frac{1}{h}(\nu_n + \delta_{\text{f}})\|Z_{t,j,h}\|_{L_2}\} \leq |\mathbb{E}Z_{t,j,h}|$. Then, under the alternative hypothesis H_1 in (4), with probability at least $1 - Ce^{-p_n} - \frac{1}{n}$ our test in (18) will reject; thus the power of the test tends to one.

Condition (i) in the theorem is satisfied as long as $|\mathbb{E}Z_{t,j,h}|$ is a large enough multiple of $\sqrt{\log(n)/n}$. In condition (ii), the signal strength $Z_{t,j,h}$ appears on both sides; however, $||Z_{t,j,h}||_{L_2}$ on the left-hand side, which admittedly is no smaller than $|\mathbb{E}Z_{t,j,h}|$ on the righthand side, is scaled by $\nu_n + \delta_f$ which we could expect to converge to zero. Thus, condition (ii) is mild as well. It is worth noting that condition (i) is itself implied by, for large enough constants $C'_{\text{alt},1}$ and $C_{\text{alt},3}$,

(a)
$$C'_{\text{alt},1}\sqrt{\log(n)/n} \leq |\int_{\Omega_h} \left[\exp\{t \, m_{0,j}(\boldsymbol{f}, x_j)\} - 1\right] d\mathbb{P}|;$$

(b) $C_{\text{alt},3}\{\frac{1}{h}\delta_{\text{f}} + r_{\text{b},m,j}\} \leq |\int_{\Omega_h} \left[\exp\{t \, m_{0,j}(\boldsymbol{f}, x_j)\} - 1\right] d\mathbb{P}|$ for $r_{\text{b},m,j}$ introduced above Theorem 3.1.

Here, condition (b) states that the effect of diversification, represented by $\delta_{\rm f}$, and the effect of smoothing the true derivative, represented by $r_{{\rm b},m,j}$, should be weak compared to the original signal. We also prove the remarks above on conditions (a) and (b) following the proof of Theorem C.6.

Proof of Theorem C.6. The original decomposition (C.20) retains its form from the proof of Theorem 3.5, but under the alternative the treatment of the various constituent terms could be different. As argued in the proof of Theorem 3.5, in the decomposition (C.20), $|\mathbb{E}Z_{t,j,h}|$ is the signal strength favoring the alternative; thus, in contrast, $I_{t,n,1}+I_{t,n,2}$ becomes the "noise" under the alternative. The sum of the left-hand sides of conditions (i) and (ii) in the statement of Theorem C.6 is, up to a multiplicative factor and with probability at least $1 - Ce^{-p_n}$, a bound on the magnitude of the noise $I_{t,n,1} + I_{t,n,2}$. The bound can be straightforwardly established by Propositions C.1 and C.2 which hold under the alternative as well. We only remark that

i) The necessity of condition (i) is due to the fact that the leading i.i.d. term within $I_{t,n,2}$ now also becomes a part of the noise that could be bounded through Hoeffding's inequality by the left-hand side of condition (i);

ii) The left-hand side of condition (ii) is similar to the left-hand side in condition (*) in Theorem 3.5 and is a bound on the term $I_{t,n,1}$ and the remainder terms in (C.22), but note that here, naturally, instead of requiring the bound to be on the order of $\mathcal{O}(n^{-1/2})$ under the null, under the alternative we require the bound to become insignificant as compared to the signal $|\mathbb{E}Z_{t,j,h}|$, which is exactly what condition (ii) requires.

Then, if the signal $|I_{t,n,3} + I_{t,n,4} + I_{t,n,5}| = |\mathbb{E}Z_{t,j,h}|$ is strong enough, as required by conditions (i) and (ii), with probability at least $1 - Ce^{-p_n}$ our test in (18) will reject, leading to the conclusion of Theorem C.6.

Next, we establish how condition (i) is implied by conditions (a) and (b). To establish condition (i), it suffices that $|I_{t,n,5}| = |\int_{\Omega_h} [\exp\{t m_{0,j}(\boldsymbol{f}, x_j)\} - 1] d\mathbb{P}|$ is strong enough, which is implied by condition (a), and moreover $|I_{t,n,5}|$ dominates $I_{t,n,3}$ and $I_{t,n,4}$. We first consider the term $I_{t,n,3}$, which arises due to the diversified projection. We first note that equivalent to (10) one can also write

$$g_{0,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j) = \int_{-1}^{1} m_{0,j}(\mathbf{H}^{\dagger}\widetilde{\boldsymbol{f}}, x_j - ah) K(a) \mathrm{d}a.$$

Thus, clearly $|g_{0,j}^{s}(\widetilde{\boldsymbol{f}}, x_{j})| \leq ||K||_{L_{1}} ||m_{0,j}||_{L_{\infty}}$ whenever $x_{j} \in \mathcal{B}_{h}$. Analogously, $|m_{0,j}^{s}(\boldsymbol{f}, x_{j})| \leq ||K||_{L_{1}} ||m_{0,j}||_{L_{\infty}}$ whenever $x_{j} \in \mathcal{B}_{h}$ as well. Now,

$$I_{t,n,3} = \tilde{\eta}_t^{s}(g_0) - \eta_t^{s}(m_0) = \int_{\Omega_h} \exp(t \,\Psi\{g_{0,j}^{s}(\widetilde{f}, x_j)\}) d\mathbb{P} - \int_{\Omega_h} \exp(t \,\{m_{0,j}^{s}(f, x_j\}) d\mathbb{P} \\ = \int_{\Omega_h} [\exp(t \,\{g_{0,j}^{s}(\widetilde{f}, x_j)\}) - \exp(t \,\{m_{0,j}^{s}(f, x_j\})] d\mathbb{P},$$

where in the transition to the last line we could afford to discard the function Ψ due to the boundedness of $g_{0,j}^{s}$ derived above. Then, by Taylor expansion and the boundedness of $m_{0,j}^{s}$ and $g_{0,j}^{s}$ derived above, followed by the Cauchy-Schwarz inequality, Lemma A.1 and then Lemma B.2,

$$|I_{t,n,3}| \leq Ct \int_{\Omega_h} |g_{0,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j) - m_{0,j}^{\mathrm{s}}(\boldsymbol{f}, x_j)| \mathrm{d}\mathbb{P} \leq Ct [\int_{\Omega_h} \{g_{0,j}^{\mathrm{s}}(\widetilde{\boldsymbol{f}}, x_j) - m_{0,j}^{\mathrm{s}}(\boldsymbol{f}, x_j)\}^2 \mathrm{d}\mathbb{P}]^{1/2} \leq C \frac{t}{h} \delta_{\mathrm{f}} \delta_{\mathrm{f}}$$

where the constant C is not dependent on t.

We next consider the term $I_{t,n,4}$. By arguments similar to those used in establishing the bound for $I_{t,n,3}$ above, but instead using the quantity $r_{b,m,j}$ introduced above Theorem 3.1 following the Cauchy-Schwarz inequality, we obtain

$$I_{t,n,4} \leq Ctr_{\mathrm{b},m,j}$$

(where the constant C again is not dependent on t).

Therefore, under condition (b), $|I_{t,n,5}|$ indeed dominates $I_{t,n,3}$ and $I_{t,n,4}$. This concludes the proof of Theorem C.6.

D Low-dimensional regime as a special case

D.1 Conditional screening in low-dimension

Our low-dimensional regime refers to the scenario where the dimension d is small. This regime can be handled as a special case of our high-dimensional regime through straightforward

modifications. First, (when d is small) in our regression model (2) we can simply set $\mathbf{F} = \mathbf{X}_{-j}$ where $\mathbf{X}_{-j} \in \mathbb{R}^{d-1}$ denotes the vector \mathbf{X} but without the *j*-th coordinate. Then, since we observe all coordinates of \mathbf{X} , the "latent" \mathbf{F} is in fact observed. Consequently, in model (2), $m_0(\mathbf{F}, X_j) = m_0(\mathbf{X}) = \mathbb{E}[Y|\mathbf{X}]$, and r = d - 1. Also assume that, in place of Assumption 3.1, we have $\mathbf{X} \in [-b, b]^d$. Then, if we wish to screen the *j*-th coordinate X_j , our conditional screening hypotheses in (3) and (4) become accordingly

$$H_0$$
: for all $\boldsymbol{x} \in [-b, b]^d$, $m_{0,j}(\boldsymbol{x}) = 0$ identically, against H_A : for some $\boldsymbol{x} \in [-b, b]^d$, we have $m_{0,j}(\boldsymbol{x}) \neq 0$.

We summarize the practical implementation of our conditional screening test in the lowdimensional regime in Algorithm D.1. The corresponding theoretical development can be obtained a special case of our theoretical results for the high-dimensional regime, by setting \tilde{F}_i to $X_{i,-j}$, the preliminary regression and tweaked regression estimators \hat{g}_n and \check{g}_n to \hat{m}_n and \check{m}_n respectively, and $\delta_f = 0$. For brevity, we omit presenting the theoretical results in the low-dimensional regime.

D.2 Simulation setup for the illustrative example

In this section we specify the data generating process for our illustrative example in the low-dimensional regime on which Figure E.2 is based. Here, the predictors $\mathbf{X} = (X_1, X_2)^{\top}$ follow a bivariate normal distribution with marginal standard deviations $\sigma_1 = 0.4$, $\sigma_2 = 0.6$, marginal means $\mu_1 = 0$, $\mu_2 = 0$, and correlation $\rho = -0.5$, while the response Y is generated as

$$Y = \sin\left(\frac{\pi}{2\sigma_1}X_1\right) + \epsilon, \quad \text{where } \epsilon \sim \mathcal{N}(0, 0.1)$$

is drawn independently of X. We choose the sample size n = 512. To specify our deep neural network, we use L = 5 hidden layers, a common width of $k_0 = 32$ per layer, and the ReLU activation function; we employ early stopping as the only regularization technique. Then, in particular, the estimator $\hat{m}_{n,j}^{s}$ plotted in Figure E.2 is based on \hat{m}_{n} produced by (D.1).

E Figures

Input: Observed samples $Y_i, \mathbf{X}_i \in \mathbb{R}^d, i \in \{1, \dots, n\}$; testing if the j^{th} has additional contributions.

- **Result:** Preliminary and refined regression estimators $\hat{m}_n(\cdot)$ and $\check{m}_n(\cdot)$, the associated smoothed derivative estimators $\hat{m}_{n,j}^{s}(\cdot)$ and $\check{m}_{n,j}^{s}(\cdot)$, $\hat{\alpha}_n(\cdot)$, and test results.
 - 1) Estimate the regression function m_0 in (2) by the preliminary neural network estimator \hat{m}_n defined as

$$\widehat{m}_n = \operatorname{argmin}_{m \in \mathcal{F}_n(d)} \frac{1}{n} \sum_{i=1}^n \left\{ Y_i - m(\boldsymbol{X}_i) \right\}^2.$$
(D.1)

Conduct cross validation (see Remark A.4), or otherwise, to find a smoothing bandwidth h.

2) For a candidate estimator α of $\alpha_{t,n}^*$, define the loss function $\widehat{R}_n^{\text{null}}(\cdot)$ as

$$\widehat{R}_n^{\text{null}}(\alpha) = \frac{1}{n} \sum_{i=1}^n \alpha^2(\boldsymbol{X}_i) - 2\frac{1}{n} \sum_{i \in \mathcal{I}_h} \alpha_j^{\text{s}}(\boldsymbol{X}_i),$$

where numerically $\alpha_j^{\rm s}$, the smoothed derivative of α , on the right-hand side can be approximated as follows: denote $a_l = -1 - \frac{1}{2N} + \frac{l}{N}$, then

$$\frac{1}{n}\sum_{i\in\mathcal{I}_h}\alpha_j^{\rm s}(\boldsymbol{X}_i)\approx-\frac{1}{hnN}\sum_{l=1}^{2N}\dot{K}(a_l)\left\{\frac{1}{n}\sum_{i\in\mathcal{I}_h}\alpha\left(\boldsymbol{X}_{i,-j},X_{i,j}+a_lh\right)\right\}.$$

We then let the estimator $\hat{\alpha}_n$ of $\alpha_{t,n}^*$ be

$$\widehat{\alpha}_n = \operatorname{argmin}_{\alpha \in \mathcal{F}_n} \widehat{R}_n^{\operatorname{null}}(\alpha).$$

3) Return $\check{m}_n = \hat{m}_n + \hat{\delta}_t \hat{\alpha}_n$, where

$$\widehat{\delta}_{t} = \frac{1}{\frac{1}{n} \sum_{i=1}^{n} \widehat{\alpha}_{n}^{2}(\boldsymbol{X}_{i})} \frac{1}{n} \sum_{i=1}^{n} \{Y_{i} - \widehat{m}_{n}(\boldsymbol{X}_{i})\} \widehat{\alpha}_{n}(\boldsymbol{X}_{i}),$$

4) For a prescribed significance level $0 < \alpha < 1$, for the fixed t test, reject the null hypothesis in our conditional screening hypotheses if

$$\frac{n^{1/2}}{t\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)}} \left|\frac{1}{n}\sum_{i=\mathcal{I}_n} \left[\exp\{t\,\widetilde{m}_j^{\mathrm{s}}(\boldsymbol{X}_i)\}-1\right]\right| > z_{1-q/2}$$

where $z_{1-q/2}$ is the 1 - q/2 quantile of a standard normal distribution, and $\|\widehat{\epsilon}\widehat{\alpha}_n\|_{L_2(\mathbb{P}_n)} = \left[\frac{1}{n}\sum_{i=1}^n \{\widehat{\epsilon}_i\widehat{\alpha}_n(\boldsymbol{X}_i)\}^2\right]^{1/2}$ is the estimated standard deviation. The uniform over t test statistics and the accompanying decision rules are analogous to those in Algorithm 2.1 for the high-dimensional regime, and are omitted.

Algorithm D.1: Partial derivative estimation & conditional screening test by deep neural network in the low dimensional regime



Figure E.1: Iterated triangular functions ζ_L in Section A.4 for L = 1, 3, 5. As L increases, the L_2 norm of the function ζ_L itself decays exponentially fast to zero, while the L_2 norm of the derivative function $\dot{\zeta}_L$ remains a constant.

F Additional application results

References

- Adamczak, R. (2008). A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electronic Journal of Probability*, 13:1000 1034.
- Andrews, D. W. K. (1994). Empirical process methods in econometrics. volume 4 of *Handbook of Econometrics*, pages 2247–2294. Elsevier.
- Anthony, M. and Bartlett, P. L. (1999). Neural Network Learning: Theoretical Foundations. Cambridge University Press.
- Bartlett, P. L., Harvey, N., Liaw, C., and Mehrabian, A. (2019). Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *Journal of Machine Learning Research*, 20(1):2285–2301.
- Bryc, W. and Smolenski, W. (1992). On the stability problem for conditional expectation. Statistics & Probability Letters, 15(1):41–46.
- Chen, X., Liao, Y., and Wang, W. (2022). Inference on time series nonparametric conditional moment restrictions using general sieves. arXiv preprint arXiv:2301.00092.
- Chen, X., Liao, Z., and Sun, Y. (2014). Sieve inference on possibly misspecified seminonparametric time series models. *Journal of Econometrics*, 178:639–658.



Figure E.2: Plot of the smoothed derivative estimator $\hat{m}_{n,j}^{s}$ (dashed black line) and the true derivative $m_{0,j}$ (solid black line), and the corresponding Root Mean Square Error (RMSE, indicated by the discrepancy between the two lines), across a range of bandwidths h. The experiment is conducted under our illustrative low-dimensional example described in Section D.2. In essence, the predictors $\mathbf{X} = (X_1, X_2)^{\top}$ follow a bivariate normal distribution, and the response Y represents a noisy observation of a sine function in X_1 . The plot shows the true and estimated derivatives when plotting Y against X_1 , and discards the variations in X_2 . To be precise, the RMSE is only calculated in the interior, indicated by the dashed gray line plotted over the dashed black line, that excludes an interval of length h from the maximum and minimum samples of X_1 . As depicted in the figures, selecting $h \approx 0.2$ achieves the optimal RMSE, striking a balance between under- and oversmoothing.



Figure F.1: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable UEMP15T26 (Part 1). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.2: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable UEMP15T26 (Part 2). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.3: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable UEMP15T26 (Part 3). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.4: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable UEMP15T26 (Part 4). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.5: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB3SMFFM (Part 1). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.6: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB3SMFFM (Part 2). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.7: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB3SMFFM (Part 3). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.8: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB3SMFFM (Part 4). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.9: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB6SMFFM (Part 1). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.10: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB6SMFFM (Part 2). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.11: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB6SMFFM (Part 3). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.



Figure F.12: FRED-MD dataset: Significance of variables after neural factor regression at the 5% significance level for the response variable TB6SMFFM (Part 4). Intervals covering one correspond to the significant idiosyncratic variables. Please refer to Figure 1 for the description of the intervals.

- Conway, J. B. (1990). A Course in Functional Analysis. Graduate Texts in Mathematics; 96. Springer Science+Business Media, New York, 2nd edition.
- Fan, J. and Gu, Y. (2023). Factor augmented sparse throughput deep ReLU neural networks for high dimensional regression. *Journal of the American Statistical Association*, pages 1–15.
- Fan, J., Gu, Y., and Zhou, W.-X. (2022). How do noise tails impact on deep ReLU networks? arXiv preprint arXiv:2203.10418.
- Fan, J. and Liao, Y. (2022). Learning latent factors from diversified projections and its applications to over-estimated and weak factors. *Journal of the American Statistical* Association, 117(538):909–924.
- Gasser, T. and Müller, H.-G. (1984). Estimating regression functions and their derivatives by the kernel method. *Scandinavian Journal of Statistics*, pages 171–185.
- Giné, E., Koltchinskii, V., and Wellner, J. A. (2003). Ratio limit theorems for empirical processes. In Giné, E., Houdré, C., and Nualart, D., editors, *Stochastic Inequalities and Applications*, pages 249–278. Birkhäuser.
- Giné, E. and Mason, D. M. (2007). On local U-statistic processes and the estimation of densities of functions of several sample variables. The Annals of Statistics, 35(3):1105– 1145.
- Giné, E. and Nickl, R. (2016). Mathematical Foundations of Infinite-Dimensional Statistical Models. Cambridge University Press.
- Gühring, I., Kutyniok, G., and Petersen, P. (2020). Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. Analysis and Applications, 18(05):803–859.
- Hall, P. and Marron, J. S. (1987). Estimation of integrated squared density derivatives. Statistics & Probability Letters, 6(2):109–115.
- Härdle, W., Marron, J. S., and Wand, M. P. (1990). Bandwidth choice for density derivatives. Journal of the Royal Statistical Society Series B: Statistical Methodology, 52(1):223–232.
- Kohler, M. and Langer, S. (2021). On the rate of convergence of fully connected deep neural network regression estimates. *The Annals of Statistics*, 49(4):2231 2249.
- Leoni, G. (2017). A First Course in Sobolev Spaces. American Mathematical Society, 2nd edition.
- Li, Q. and Racine, J. S. (2006). Nonparametric Econometrics: Theory and Practice. Princeton University Press.
- Protter, M. H. and Morrey, Jr., C. B. (2012). *Intermediate Calculus*. Springer Science & Business Media.
- Resnick, S. (2005). A Probability Path. Springer.

- Schmidt-Hieber, J. (2020). Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4):1875–1897.
- Stein, E. M. and Shakarchi, R. (2009). *Real Analysis: Measure Theory, Integration, and Hilbert spaces.* Princeton University Press.
- Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *The* Annals of Statistics, pages 1040–1053.
- Telgarsky, M. (2015). Representation benefits of deep feedforward networks. arXiv preprint arXiv:1509.08101.
- van der Vaart, A. W. and Wellner, J. A. (1996). Weak Convergence and Empirical Processes: With Applications to Statistics. Springer.
- Yarotsky, D. (2017). Error bounds for approximations with deep ReLU networks. Neural Networks, 94:103–114.