

Testing Mediation Effects Using Logic of Boolean Matrices

Chengchun Shi and Lexin Li

*London School of Economics and Political Science
and University of California at Berkeley*

Abstract

A central question in high-dimensional mediation analysis is to infer the significance of individual mediators. The main challenge is that the total number of potential paths that go through any mediator is super-exponential in the number of mediators. Most existing mediation inference solutions either explicitly impose that the mediators are conditionally independent given the exposure, or ignore any potential directed paths among the mediators. In this article, we propose a novel hypothesis testing procedure to evaluate individual mediation effects, while taking into account potential interactions among the mediators. Our proposal thus fills a crucial gap, and greatly extends the scope of existing mediation tests. Our key idea is to construct the test statistic using the logic of Boolean matrices, which enables us to establish the proper limiting distribution under the null hypothesis. We further employ screening, data splitting, and decorrelated estimation to reduce the bias and increase the power of the test. We show that our test can control both the size and false discovery rate asymptotically, and the power of the test approaches one, while allowing the number of mediators to diverge to infinity with the sample size. We demonstrate the efficacy of the method through simulations and a neuroimaging study of Alzheimer’s disease. A Python implementation of the proposed procedure is available at <https://github.com/callmespring/LOGAN>.

Keywords: Boolean matrix; Directed acyclic graph; Gaussian graphical model; High-dimensional inference; Mediation analysis; Neuroimaging analysis.

1 Introduction

Mediation analysis is an important tool in scientific studies. It seeks to identify and explain the mechanism, or pathway, that underlies an observed relationship between an exposure and an outcome variable, through the inclusion of an intermediary variable, known as a mediator. It decomposes the effect of exposure on the outcome into a direct effect and an indirect effect, the latter of which is often of primary interest and has important intervention consequences (Pearl, 2001). Mediation analysis was first proposed with a single mediator in social science (Baron and Kenny, 1986). In recent years, it is receiving increasing attention, and has been extended to the settings of multivariate and high-dimensional mediators. It is now widely used in a large variety of scientific applications, including psychology (MacKinnon and Fairchild, 2009), genomics (Huang and Pan, 2016), genetic epidemiology (Huang, 2018), and neuroscience (Zhao and Luo, 2016).

In mediation analysis with high-dimensional mediators, a fundamental but challenging question is how to infer the significance of individual mediators. The main difficulty is the sheer number of possible paths that go through all combinations of the mediators. Consequently, the total number of potential paths that go through any mediator is super-exponential in the number of mediators, rendering almost any existing testing procedure ineffective. To circumvent this issue, most existing mediation inference solutions either explicitly impose that the mediators are conditionally independent given the exposure, or simply ignore any potential directed paths among the mediators. Such simplifications substantially reduce the complexity of the hypotheses to test. Adopting this conditional independence assumption, Boca et al. (2014) proposed a permutation test with family-wise error control, while Zhang et al. (2016) proposed a screening-and-testing assisted approach. Huang and Pan (2016) proposed a transformation model and assumed conditional independence for the transformed mediators. Sampson et al. (2018) and Djordjilović et al. (2019) directly tested whether each mediator is independent of the exposure or conditionally independent of the outcome given the exposure, ignoring mediator-by-mediator interactions, while controlling for family-wise error rate or false discovery rate in multiple testing. Whereas these tests have been demonstrated effective in numerous applications, they all ignored potential paths and interactions among the mediators. Even though this strategy may be plausible in some applications, it may not hold true in others. For instance, in our brain imaging mediation

analysis study in Section 7, different brain regions are conceived to influence each other. In genetics studies, different genes are expected to interact with each other (Chakraborty et al., 2018). Actually, such examples are often the rule rather than the exception. Therefore, it is of great importance to develop a mediation testing method that takes directed paths and interactions among the mediators into consideration.

Recently, Chakraborty et al. (2018) made an important step forward for inference of mediation effects while allowing mediator interactions. They formulated the structure of the exposure, the potential mediators, and the outcome as a directed acyclic graph (DAG). They defined the individual mediation effect of a given mediator as the summation of all the effects of the exposure on the outcome that can be attributed to that mediator. They then established the corresponding confidence interval for their interventional calculus type estimator of the mediation effect. However, the effects along different paths may cancel each other, resulting in a zero individual mediation effect in the summation. Rather than taking average and cancelling out the total effect, we argue this type of mediator is important and should be identified by the inferential test too. See Section 2.2 for more discussion.

There have also been some recent proposals of penalized sparse estimation of mediation effects (Zhao and Luo, 2016; Nandy et al., 2017). In addition, there is a large body of literature studying penalized estimation of directed acyclic graph given observational data (see, e.g., van de Geer and Bühlmann, 2013; Zheng et al., 2018; Yuan et al., 2019, and the references therein). However, estimation is an ultimately different problem from inference. Although both can in effect identify important mediators or links, estimation does not produce an explicit quantification of statistical significance, and does not explicitly control the false discovery. As such, we are targeting a completely different problem than those estimation approaches. More recently, Li et al. (2019) developed a constrained likelihood ratio test to infer individual links or some given directed paths of a DAG. Nevertheless, their hypotheses are very different from our problem of inferring significant mediators.

In this article, we propose a novel hypothesis testing procedure to evaluate individual mediation effects, which takes into account directed paths among the mediators and is equipped with statistical guarantees. A key ingredient of our proposal is to construct the test statistic using the logic of Boolean matrices, which allows us to establish the proper limiting distribution under the null hypothesis. In comparison, the asymptotic properties of the test statistic built on the usual matrix operations are extremely challenging to es-

tablish. In addition, the Boolean matrix-based test statistic can be naturally coupled with a screening procedure. This helps scale down the number of potential paths to a moderate level, and in turn reduces the variance of the test statistic, and enhances the power of the test considerably. Furthermore, we use a data splitting strategy to ensure a valid type-I error rate control for our test under minimal conditions on the screening. We employ some state-of-the-art estimator of DAG (Zheng et al., 2018) to form an initial estimator of the directed paths. We also devise a decorrelated estimator to reduce potential bias induced by high-dimensional mediators. Consequently, it ensures the resulting estimator is \sqrt{n} -consistent and asymptotically normal. We then employ a multiplier bootstrap method to obtain accurate critical values. Finally, we couple our test for the significance of an individual mediator with a multiple testing procedure (Djordjilović et al., 2019) to control the false discovery rate (FDR) of simultaneous testing of multivariate mediators.

Our contributions are multi-fold. Scientifically, rigorous inference of mediation effects is a vital and long-standing problem. But nearly all existing solutions ignore potential interactions among the mediators. Our proposal thus fills a crucial gap, extends the scope of existing tests, and offers a useful inferential tool to a wide range of scientific applications. Methodologically, our proposal integrates the logic of Boolean matrices, DAG estimation, screening, data splitting, and decorrelated estimation to reduce the bias and increase the power of the test. It is ultimately different from the test of Chakraborty et al. (2018), which defined the mediation effect through averaging, required the DAG selection consistency, focused on dealing with the equivalence class of DAG estimators, and did not consider multiple testing. By contrast, our method targets a different, and in our opinion, a more general definition of mediation effect, does not require the DAG selection consistency, and mostly focuses on the single DAG situation. We discuss the extension of the test to the equivalence class situation in Section 8. We also compare our test with that of Chakraborty et al. (2018) numerically, and show our method is empirically more powerful while achieving a valid type-I error control. Theoretically, we systematically study the asymptotic properties of our test, while allowing the number of mediators to diverge to infinity with the sample size. We show that our test can control both the size and FDR asymptotically, and the power of the test approaches one. As a by-product, we derive an oracle inequality for the estimated DAG by the method of Zheng et al. (2018), which is needed to establish the consistency of our test, but is not available in Zheng et al. (2018).

The rest of the article is organized as follows. We define our hypotheses in Section 2, and develop the test statistics based on the logic of Boolean matrices in Section 3. We propose the testing procedures in Section 4, and investigate the asymptotic properties in Section 5. We present the simulations in Section 6, and a neuroimaging application in Section 7. We conclude the paper in Section 8, and relegate all proofs to the supplementary appendix.

2 Hypotheses

In this section, we first present the Gaussian graphical model, based on which we formulate our mediation testing problem. We then formally develop the hypotheses we aim to test, and compare with the alternative formulation in Chakraborty et al. (2018).

2.1 Gaussian graphical model

Consider an exposure variable E , a set of potential mediators $\mathbf{M} = (M_1, \dots, M_d)^\top$, and an outcome variable Y . Let $\mathbf{X} = (E, \mathbf{M}^\top, Y)^\top$ collect all the variables, and assume \mathbf{X} follows the linear structural equation model,

$$\mathbf{X} - \boldsymbol{\mu}_0 = \mathbf{W}_0(\mathbf{X} - \boldsymbol{\mu}_0) + \boldsymbol{\varepsilon}, \quad (1)$$

where $\boldsymbol{\mu}_0 = \mathbb{E}(\mathbf{X})$, \mathbf{W}_0 is the $(d+2) \times (d+2)$ coefficient matrix, and $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \dots, \varepsilon_{d+1})^\top$ is the mean-zero vector of errors. The matrix \mathbf{W}_0 specifies the directional relationships among the variables in \mathbf{X} , which can be encoded by a directed graph. Let X_j denote the $(j+1)$ th element of \mathbf{X} , $j = 0, \dots, d+1$. For $i, j \in \{0, 1, \dots, d+1\}$, if X_i is a direct cause of X_j , then an arrow is drawn from X_i to X_j , i.e., $X_i \rightarrow X_j$, and $W_{0,j,i} \neq 0$. In this case, X_i is called a parent of X_j , and X_j a child of X_i . For an integer $k \geq 1$, a k -step directed path between X_i and X_j is a sequence of distinct nodes from X_i to X_j : $X_i \rightarrow X_{i_1} \rightarrow \dots \rightarrow X_{i_{k-1}} \rightarrow X_j$, for some $\{i_k\}_{1 \leq l < k}$. In this case, X_i is called an ancestor of X_j , and X_j a descendant of X_i . For model (1) and the associated directed graph, we impose a set of conditions. Specifically,

- (A1) The directed graph is acyclic; i.e., no variable is an ancestor of itself.
- (A2) No potential mediator M_i is a direct cause of the exposure E , and the outcome Y is not a direct cause of neither the exposure E nor any mediator M_i , $i = 1, \dots, d$.
- (A3) The errors ε_i , $i = 0, 1, \dots, d+1$, are jointly normally distributed and independent.

In addition, the error variances $\sigma_i^2 = \text{Var}(\varepsilon_i)$, $i = 0, \dots, d+1$, are constant; i.e., $\sigma_0^2 = \sigma_1^2 = \dots = \sigma_d^2 = \sigma_{d+1}^2 = \sigma_*^2$ for some constant $\sigma_* > 0$.

These model assumptions are generally mild, and are often imposed in the DAG and mediation analysis literature. Specifically, Condition (A1) implies that \mathbf{W}_0 is a lower-triangular matrix, up to a permutation of the rows and columns. Condition (A2) implies that the first row of \mathbf{W}_0 and the last column of \mathbf{W}_0 are both zero vectors. Condition (A3) basically specifies that \mathbf{X} follows a Gaussian graphical model. By Gram-Schmidt orthogonalization, any Gaussian DAG model can always be represented by (1) with independent errors. In addition, the constant variance condition in (A3) ensures that, under the Gaussian graphical model (1), \mathbf{W}_0 is identifiable (Peters and Bühlmann, 2014, Theorem 1). This avoids the situation of the equivalence class of DAG, and a similar condition has been adopted in Yuan et al. (2019) as well. We note that it is possible to relax (A3) by requiring $\sigma_1^2 = \dots = \sigma_d^2$; i.e., excluding the variance requirement on the exposure and the outcome. We discuss this relaxation in more details in Section .2 of the appendix. Moreover, we also discuss the extension of our method to the unequal variance case in Section 8.

2.2 Mediation effects and hypotheses

For a directed path $\zeta : E \rightarrow M_{i_1} \rightarrow \dots \rightarrow M_{i_k} \rightarrow Y$ for some $\{i_t\}_{1 \leq t \leq k} \subseteq \{1, \dots, d\}$, we define the total effect of E on Y attributed to this path as,

$$\omega_\zeta = W_{0,i_1,0} \left(\prod_{t=0}^{k-1} W_{0,i_{t+1},i_t} \right) W_{0,d+1,i_k}, \quad (2)$$

where $W_{0,i,j}$ is the (i, j) th entry of \mathbf{W}_0 . If such a path does not exist, we have $\omega_\zeta = 0$. This definition of total effect ω_ζ plays a central role in our definition of mediation effect.

Based on (2), we formally state our hypotheses regarding the significance of an individual mediator M_q , for an integer $q = 1, \dots, d$,

$$\begin{aligned} H_0(q) : \omega_\zeta &= 0, \quad \text{for all } \zeta \text{ that passes through } M_q, & \text{versus} \\ H_1(q) : \omega_\zeta &\neq 0, \quad \text{for some } \zeta \text{ that passes through } M_q. \end{aligned} \quad (3)$$

When the alternative hypothesis in (3) holds, we call M_q a significant mediator.

We observe that, the hypotheses in (3) can be reformulated as the following equivalent pair of hypotheses. That is, for any integer $j = 1, \dots, d+1$, let $\text{ACT}(j, \mathbf{W}_0)$ denote the set

of true ancestors of X_j ; i.e., $i \in \text{ACT}(j, \mathbf{W}_0)$ if and only if X_i is an ancestor of X_j . Then the pair of hypotheses (3) is equivalent to the following pair of hypotheses,

$$\begin{aligned} H_0(q) : 0 \notin \text{ACT}(q, \mathbf{W}_0) \quad \text{or} \quad q \notin \text{ACT}(d+1, \mathbf{W}_0), \quad \text{versus} \\ H_1(q) : 0 \in \text{ACT}(q, \mathbf{W}_0) \quad \text{and} \quad q \in \text{ACT}(d+1, \mathbf{W}_0). \end{aligned} \quad (4)$$

Next, we consider a pair of hypotheses that lead to (4). For any $q_1 = 0, \dots, d$, $q_2 = 1, \dots, d+1$, we consider the following pair of hypotheses,

$$H_0(q_1, q_2) : q_1 \notin \text{ACT}(q_2, \mathbf{W}_0), \quad \text{versus} \quad H_1(q_1, q_2) : q_1 \in \text{ACT}(q_2, \mathbf{W}_0). \quad (5)$$

We observe that, the null hypothesis $H_0(q)$ in (4) can be decomposed into a union of the two null hypotheses $H_0(0, q)$ and $H_0(q, d+1)$ that are defined in (5). Suppose $p(q_1, q_2)$ is a valid p -value for $H_0(q_1, q_2)$ in (5). According to the union-intersection principle, $\max\{p(0, q), p(q, d+1)\}$ is a valid p -value for testing $H_0(q)$ in (4). Therefore, we aim at (5) in the subsequent development of our testing procedure.

We have defined a significant mediator through (3). There is an alternative definition employed by Chakraborty et al. (2018). Specifically, they considered the hypotheses,

$$H_0^*(q) : \sum \omega_\zeta = 0, \quad \text{versus} \quad H_1^*(q) : \sum \omega_\zeta \neq 0, \quad (6)$$

where the summation is taken for all ζ that pass through M_q . Chakraborty et al. (2018) called M_q a significant mediator when the alternative hypothesis in (6) holds. We, however, prefer our definition of a significant mediator that is built on (3) instead of (6). This is because the effects along the path ζ may cancel out with each other, resulting in a zero sum, even though there are significant positive and negative mediation effects along ζ . As an illustration, we devise a simple example as shown in Figure 1(a). For the mediator M_2 , two paths, $E \rightarrow M_2 \rightarrow Y$ and $E \rightarrow M_2 \rightarrow M_3 \rightarrow Y$, both pass through M_2 . The aggregated total effect following (6) is $\sum_\zeta \omega_\zeta = 1 \times \{-1 + (-1) \times (-1)\} = 0$. Similarly, we can show the aggregated total effect of M_3 is zero as well. However, both M_2 and M_3 have positive and negative mediation effects, and should be viewed as significant mediators.

We conclude this section by computing the explicit number of potential paths that go through any mediator M_q . For an integer $k = 2, 3, \dots, d+1$, the total number of k -step potential paths that go through M_q , by the combinatorial theory, is $N_k(q) = \binom{d-1}{k-2}(k-1)!$. Then the total number of potential paths that go through M_q is $N(q) = \sum_{k=2}^{d+1} N_k(q) \geq$




Figure 1: Left panel: a DAG with five nodes, where node 0 is the exposure variable, node 4 is the outcome variable, and nodes 1 to 3 are the mediators. Right panel: a DAG with three nodes.

$\mathcal{N}_{d+1}(q) = d!$. As a result, it is highly nontrivial to test the significance of an individual mediator if we take into account all the potential paths among the mediators.

3 Test Statistics

In this section, we first consider a potential test statistic built on the power of an estimator of the coefficient matrix \mathbf{W}_0 in model (1), and discuss its limitation. We then present our main idea, the logic of Boolean matrices, and the test statistic built on it.

3.1 Power of matrices

Matrices and vectors in this paper start the index from zero. For any matrix \mathbf{A} , let $|\mathbf{A}|$ denote the matrix of the same dimension whose (i, j) th entry is $|A_{i,j}|$. We first connect the null hypothesis $H_0(q_1, q_2)$ in (5) with the coefficient matrix \mathbf{W}_0 in model (1). Recall that $H_0(q_1, q_2)$ means X_{q_1} is not an ancestor of X_{q_2} . We have the next lemma.

Lemma 1. *The null $H_0(q_1, q_2)$ holds if and only if $(|\mathbf{W}_0|^k)_{q_2, q_1} = 0$ for any $k = 1, \dots, d$.*

We sketch the proof of this lemma here, which is to facilitate our understanding of the problem. The key observation is that, the (q_2, q_1) th entry of $|\mathbf{W}_0|^k$ is the sum of the absolute values of the total effects along all k -step paths from X_{q_1} to X_{q_2} . For instance, for $k = 2$,

$$(|\mathbf{W}_0|^2)_{q_2, q_1} = \sum_{j=0}^{d+1} |W_{0,j,q_1}| |W_{0,q_2,j}| = \sum_{j=1}^d |W_{0,j,q_1}| |W_{0,q_2,j}|,$$

where the last equality is due to that the first row and last column of \mathbf{W}_0 are zero vectors because of Condition (A2). If there exists a two-step path from X_{q_1} to X_{q_2} , by definition,

$W_{0,j,q_1} \neq 0$ and $W_{0,q_2,j} \neq 0$ for some $j = 1, \dots, d$, which is equivalent to $(|\mathbf{W}_0|^2)_{q_2,q_1} \neq 0$. Similarly, there exists a k -step path from X_0 to M_q if and only if $(|\mathbf{W}_0|^k)_{q_2,q_1} \neq 0$.

Let $\widehat{\mathbf{W}}$ be some consistent estimator for \mathbf{W}_0 . In view of Lemma 1, it is natural to construct a test statistic based on $\{(|\widehat{\mathbf{W}}|^k)_{q_2,q_1}\}_{1 \leq k \leq d}$. The major difficulty with this potential test statistic, however, is that it is unclear whether $(|\widehat{\mathbf{W}}|^k)_{q_2,q_1}$ has a well tabulated limiting distribution under $H_0(q_1, q_2)$. To better illustrate this, we first consider the case when $k = 2$. We have $(|\widehat{\mathbf{W}}|^2)_{q_2,q_1} = \sum_{j=1}^d |\widehat{W}_{j,q_1}| |\widehat{W}_{q_2,j}|$. Under $H_0(q_1, q_2)$ and the acyclic constraint (A1), for any j , either W_{0,j,q_1} or $W_{0,q_2,j}$ equals zero. Suppose each \widehat{W}_{q_1,q_2} is root- n consistent to W_{0,q_1,q_2} , and the mediator dimension d is fixed. Then we can show $\sqrt{n}(|\widehat{\mathbf{W}}|^2)_{q_2,q_1}$ is asymptotically equivalent to $\sum_{1 \leq j \leq d} \sqrt{n} (|\widehat{W}_{j,q_1} - W_{0,j,q_1}| |W_{0,q_2,j}| + |W_{0,j,q_1}| |\widehat{W}_{q_2,j} - W_{0,q_2,j}|)$. The limiting distribution, however, is not well-studied even in the fixed- d scenario. When k is large, or when the mediator dimension d diverges with the sample size n , the derivation of the asymptotic property of $(|\widehat{\mathbf{W}}|^k)_{q_2,q_1}$ becomes more complicated due to the addition and multiplication operations involved in $(|\widehat{\mathbf{W}}|^k)_{q_2,q_1}$. Therefore, the test statistic based on $|\widehat{\mathbf{W}}|^k$ may not be suitable for our purpose of testing significant mediators.

3.2 Logic of Boolean matrices

To overcome the difficulty regarding $|\widehat{\mathbf{W}}|^k$, and motivated by the logic of Boolean matrices, we define a new matrix multiplication operator and a new matrix addition operator to replace the usual matrix multiplication and addition operations. Specifically, for any two real-valued matrices $\mathbf{A}_1 = \{a_{1,i,j}\}_{ij} \in \mathbb{R}^{q_1 \times q_2}$, $\mathbf{A}_2 = \{a_{2,i,j}\}_{ij} \in \mathbb{R}^{q_2 \times q_3}$, we define $\mathbf{A}_1 \otimes \mathbf{A}_2$ to be a $q_1 \times q_3$ matrix whose (i, j) th entry equals $\max_{k \in \{1, \dots, q_2\}} \min(a_{1,i,k}, a_{2,k,j})$. That is, we replace the multiplication operation in the usual matrix multiplication with the minimum operator, and replace the addition operation with the maximum operator. When $\mathbf{A}_1, \mathbf{A}_2$ are binary matrices, the minimum and maximum operators are equivalent to the logic operators “and” and “or” in Boolean algebra. The defined “ \otimes ” operator is then equivalent to the Boolean matrix multiplication operator. Moreover, for any two real-valued matrices $\mathbf{A}_1 = \{a_{1,i,j}\}_{ij}, \mathbf{A}_2 = \{a_{2,i,j}\}_{ij} \in \mathbb{R}^{q_1 \times q_2}$, we define $\mathbf{A}_1 \oplus \mathbf{A}_2$ to be a $q_1 \times q_2$ matrix whose (i, j) th entry equals $\max(a_{1,i,j}, a_{2,i,j})$. When $\mathbf{A}_1, \mathbf{A}_2$ are binary matrices, the defined “ \oplus ” operator is equivalent to the Boolean matrix addition operator.

Given the new definition of the multiplication and addition operators, we define $|\mathbf{W}|_0^{(k)} =$

$|\mathbf{W}_0|^{(k-1)} \otimes |\mathbf{W}_0|$ in a recursive fashion, for any $k \geq 1$. Next, we connect the null hypothesis $H_0(q_1, q_2)$ in (5) with the newly defined $|\mathbf{W}_0|^{(k)}$. Its proof is given in the appendix.

Lemma 2. *The null $H_0(q_1, q_2)$ holds if and only if $(|\mathbf{W}_0|^{(k)})_{q_2, q_1} = 0$ for any $k = 1, \dots, d$.*

Aggregating $|\mathbf{W}_0|^{(k)}$ for all k -step paths, $k = 1, \dots, d$, leads to the following definition,

$$\mathbf{W}_0^* = |\mathbf{W}_0| \oplus |\mathbf{W}_0|^{(2)} \oplus \dots \oplus |\mathbf{W}_0|^{(d)}.$$

We next define two matrices \mathbf{B}_0 and \mathbf{B}_0^* that are the binary versions of \mathbf{W}_0 and \mathbf{W}_0^* ,

$$(\mathbf{B}_0)_{i,j} = \begin{cases} 1, & \text{if } W_{0,i,j} \neq 0, \\ 0, & \text{otherwise,} \end{cases} \quad \text{and} \quad \mathbf{B}_0^* = \mathbf{B}_0 \oplus \mathbf{B}_0^{(2)} \oplus \dots \oplus \mathbf{B}_0^{(d)}.$$

Then Lemma 2 immediately implies the next result.

Corollary 1. *The null $H_0(q_1, q_2)$ holds if and only if $(\mathbf{W}_0^*)_{q_2, q_1} = 0$ and $(\mathbf{B}_0^*)_{q_2, q_1} = 0$.*

Corollary 1 suggests some natural test statistic for our hypotheses. Again, let $\widehat{\mathbf{W}}$ be some consistent estimator for \mathbf{W}_0 , and let $\widehat{\mathbf{W}}^* = |\widehat{\mathbf{W}}| \oplus |\widehat{\mathbf{W}}|^2 \oplus \dots \oplus |\widehat{\mathbf{W}}|^d$. We further define a thresholded binary version $\widehat{\mathbf{B}}(c)$ and $\widehat{\mathbf{B}}^*(c)$, for a given thresholding value c , as,

$$\{\widehat{\mathbf{B}}(c)\}_{i,j} = \begin{cases} 1, & \text{if } |\widehat{W}_{i,j}| > c, \\ 0, & \text{otherwise,} \end{cases} \quad \text{and} \quad \widehat{\mathbf{B}}^*(c) = \widehat{\mathbf{B}}(c) \oplus \widehat{\mathbf{B}}^{(2)}(c) \oplus \dots \oplus \widehat{\mathbf{B}}^{(d)}(c). \quad (7)$$

In view of Corollary 1, we expect $\widehat{\mathbf{W}}_{q_2, q_1}^*$ to be small under $H_0(q_1, q_2)$, and we reject $H_0(q_1, q_2)$ when $(\widehat{\mathbf{W}}^*)_{q_2, q_1} > c$ for some thresholding value c , or equivalently, when $\{\widehat{\mathbf{B}}^*(c)\}_{q_2, q_1} = 1$. We then build a test statistic based on $\widehat{\mathbf{W}}^*$.

Unlike the usual power of the matrix $|\widehat{\mathbf{W}}|^k$, the limiting distribution of $\widehat{\mathbf{W}}^*$ based on the logic of Boolean matrices is more tractable. Specifically, under the null hypothesis $H_0(q_1, q_2)$, for any potential path $X_{q_1} \rightarrow X_{j_1} \rightarrow \dots \rightarrow X_{j_k} \rightarrow X_{q_2}$, such that

$$\widehat{W}_{j_1, q_1} \neq 0, \min_{t \in \{1, \dots, k-1\}} |\widehat{W}_{j_{t+1}, j_t}| \neq 0, \quad \text{and} \quad \widehat{W}_{q_2, j_k} \neq 0, \quad (8)$$

there exist some distinct integers $\ell_1, \ell_2 \in \{q_1, j_1, \dots, j_k, q_2\}$ as functions of $(q_1, \{j_t\}_{1 \leq t \leq k}, q_2)$, such that $W_{0, \ell_2, \ell_1} = 0$. It then follows that,

$$\begin{aligned} \sqrt{n}(\widehat{\mathbf{W}}^*)_{q_2, q_1} &\leq \max_{\substack{k \in \{1, \dots, d\} \\ 1 \leq j_1, \dots, j_k \leq d+2}} \sqrt{n} |\widehat{W}_{\ell_2, \ell_1}| = \max_{\substack{k \in \{1, \dots, d\} \\ 0 \leq j_1, \dots, j_k \leq d+1}} \sqrt{n} |\widehat{W}_{\ell_2, \ell_1} - W_{0, \ell_2, \ell_1}| \\ &\leq \max_{\substack{k \in \{1, \dots, d\} \\ 0 \leq j_1, \dots, j_k \leq d+1}} \max_{\substack{0 \leq t \leq k \\ j_0 = q_1, j_{k+1} = q_2}} \sqrt{n} |\widehat{W}_{j_{t+1}, j_t} - W_{0, j_{t+1}, j_t}|, \end{aligned} \quad (9)$$

where the first maximum is taken over all such k and (j_1, \dots, j_k) that satisfy (8). When the nonzero entries of $\sqrt{n}(\widehat{\mathbf{W}} - \mathbf{W}_0)$ are asymptotically normal, the right-hand-side of (9) converges in distribution to a maximum of some normal random variables in absolute values. Its α th upper quantile can be consistently estimated via multiplier bootstrap. This forms the basis of our proposed testing procedure.

On the other hand, the test outlined above has some limitations. One is that this test can be conservative when \mathbf{W}_0 is highly sparse but $\widehat{\mathbf{W}}$ is not. Another limitation is that it requires the support of $\widehat{\mathbf{W}}$ to be fixed. When this fixed support condition does not hold, it would lead to an inflated type-I error rate. To address these limitations, we next develop a testing procedure that couples such a test with screening and data sample splitting to enhance its power as well as to ensure its validity.

4 Testing Procedure

In this section, we first present our full testing procedure for inference of an individual mediator. We next describe in detail some major steps of this testing procedure. Finally, we present a multiple testing procedure for simultaneous inference of multivariate mediators with a proper FDR control. Given that our test is constructed based on the LOGic of booleAN matrices, we refer our testing method as LOGAN.

4.1 The complete algorithm

Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ denote n i.i.d. copies of \mathbf{X} , generated according to model (1). Step 1 of our testing procedure is to randomly divide the observed data into two equal halves $\{\mathbf{X}_i\}_{i \in \mathcal{I}_1} \cup \{\mathbf{X}_i\}_{i \in \mathcal{I}_2}$, where \mathcal{I}_ℓ is the set of indices of subsamples, $\ell = 1, 2$. The purpose of data splitting is to ensure our test achieves a valid type-I error rate under minimal conditions when coupled with a screening step. In recent years, data splitting has been commonly used in high-dimensional estimation and inference (e.g., Chernozhukov et al., 2018; Newey and Robins, 2018; Barber and Candés, 2019; Romano and DiCiccio, 2019). One issue with data splitting is the potential loss of power due to the usage of only a fraction of data. There have been studies showing that data splitting may improve the power in some cases (Rubin et al., 2006; Romano and DiCiccio, 2019). In our setting, we construct two test statistics based on both halves of data, then combine them to derive the final decision rule. We show

the test constructed this way achieves a good power both asymptotically and numerically. Moreover, one may follow the idea of Meinshausen et al. (2009) to carry out the binary split more than once, then combine the p -values from all splits. This strategy also helps mitigate the randomness the single data splitting introduces. In the regression setting, Meinshausen et al. (2009) showed empirically that this multi-split strategy improves the power than a single-split. We develop a multi-split version of our test in Section .1 of the appendix, and show its improvement in power numerically. We also note that, one may employ the multi-split strategy of Romano and DiCiccio (2019, Section 4.2.1) for power improvement. Meanwhile, these improvements all come with a price of increased computational costs.

Step 2 is to compute an initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ for \mathbf{W}_0 , given each half of the data $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell}$, $\ell = 1, 2$. Several methods can be used here, e.g., Zheng et al. (2018); Yuan et al. (2019). We only require $\widetilde{\mathbf{W}}^{(\ell)}$ to be consistent to \mathbf{W}_0 . This requirement is considerably weaker than requiring $\widetilde{\mathbf{W}}^{(\ell)}$ to be selection consistent; i.e., $\mathbb{I}(\widetilde{W}_{i,j}^{(\ell)} = 0) = \mathbb{I}(W_{0,i,j} = 0)$ for any $i, j = 0, \dots, d+1$, where $\mathbb{I}(\cdot)$ is the indicator function. See Section 4.2 for more details.

Step 3 is to compute the binary matrix $\widehat{\mathbf{B}}^{(\ell)}$ for \mathbf{B}_0 , given the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$, using (7) with $c = 0$. This step is straightforward, and the main purpose is to allow the subsequent decorrelated estimation step to focus only on those nonzero elements in $\widehat{\mathbf{B}}^{(\ell)}$. It thus acts as a screening step, and in effect reduces the number of potential paths to a moderate level. As a benefit, it reduces the variance of the Boolean matrix-based test statistic, and increases the power of the test. See Section 4.3 for more details.

Step 4 is to compute a decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ using a cross-fitting procedure. We use one set of samples \mathcal{I}_ℓ to obtain the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ and the binary version $\widehat{\mathbf{B}}^{(\ell)}$ to screen out the zero entries, then use the other set of samples \mathcal{I}_ℓ^c to compute the entries of the decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$. This decorrelated estimation step is to reduce the bias of $\widetilde{\mathbf{W}}^{(\ell)}$ under the setting of high-dimensional mediators. Moreover, it guarantees the entry of $\widehat{\mathbf{W}}^{(\ell)}$, $\widehat{W}_{j_1, j_2}^{(\ell)}$, is \sqrt{n} -consistent and asymptotically normal. It adopts the debiasing idea that is commonly used for statistical inference of low-dimensional parameters in high-dimensional models (Zhang and Zhang, 2014; Ning and Liu, 2017). See Section 4.3 for more details.

Step 5 is to use a bootstrap-based procedure to compute the critical values. Let $\widehat{\mathbf{W}}^{*(\ell)} = |\widehat{\mathbf{W}}^{(\ell)}| \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(2)} \oplus \dots \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(d)}$. Similar to (9), we have

$$\sqrt{|\mathcal{I}_\ell^c|} (\widehat{\mathbf{W}}^{*(\ell)})_{q_2, q_1} \leq \max_{\substack{k \in \{1, \dots, d\} \\ 0 \leq j_1, \dots, j_k \leq d+1}} \max_{\substack{0 \leq t \leq k \\ j_0 = q_1, j_{k+1} = q_2}} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{j_{t+1}, j_t}^{(\ell)} - W_{0, j_{t+1}, j_t}^{(\ell)}|.$$

Algorithm 1 Testing procedure for inference of an individual mediator.

Input: The data $\mathbf{X}_1, \dots, \mathbf{X}_n$, $1 \leq q \leq d$, and the significance level $0 < \alpha < 1$.

Step 1. Randomly divide $\{1, 2, \dots, n\}$ into two disjoint subsets $\mathcal{I}_1 \cup \mathcal{I}_2$ of equal sizes.

Step 2. Compute an initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ for \mathbf{W}_0 , given $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell}$, $\ell = 1, 2$.

Step 3. Compute the binary estimator $\widehat{\mathbf{B}}^{(\ell)}$ for \mathbf{B}_0 , given $\widetilde{\mathbf{W}}^{(\ell)}$, $\ell = 1, 2$, which is to be used for screening and also ancestor estimation in the next step.

Step 4. Compute the decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ for \mathbf{W}_0 , given $\widetilde{\mathbf{W}}^{(\ell)}$, $\widehat{\mathbf{B}}^{(\ell)}$ and $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell}$, $\ell = 1, 2$.

(4a) Estimate the ancestors of M_q , for $q = 1, \dots, d+1$, by $\text{ACT}(q, \widetilde{\mathbf{W}}^{(\ell)}) = \left\{ 1 \leq j \leq d : \{\widehat{\mathbf{B}}^{*(\ell)}\}_{q,j} \neq 0 \right\}$, where $\widehat{\mathbf{B}}^{*(\ell)} = |\widehat{\mathbf{B}}^{(\ell)}| \oplus |\widehat{\mathbf{B}}^{(\ell)}|^{(2)} \oplus \dots \oplus |\widehat{\mathbf{B}}^{(\ell)}|^{(d)}$.

(4b) Update the j th row of $\widetilde{\mathbf{W}}^{(\ell)}$, for $1 \leq j \leq d$, by fitting a penalized regression with $\{X_{i,j}\}_{i \in \mathcal{I}_\ell}$ being the response and $\{X_{i,k}\}_{i \in \mathcal{I}_\ell, \widetilde{W}_{j,k}^{(\ell)} \neq 0}$ being the predictors. Denote the updated estimator as $\overline{\mathbf{W}}^{(\ell)}$.

(4c) Compute the decorrelated estimator $\widehat{W}_{j_1, j_2}^{(\ell)}$, for any (j_1, j_2) such that $\widehat{B}_{j_1, j_2}^{(\ell)} \neq 0$, given $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell^c}$, $\text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})$, $\widehat{\mathbf{B}}^{(\ell)}$, and $\overline{\mathbf{W}}^{(\ell)}$.

Step 5. Compute the critical values using the bootstrap procedure, given $\widehat{\mathbf{W}}^{(\ell)}$, $\widehat{\mathbf{B}}^{(\ell)}$, and $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell}$, $\ell = 1, 2$.

(5a) Compute the critical value $\widehat{c}^{(\ell)}(0, q)$ of $\max_{(i,j) \in \mathcal{S}(0,q,\widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{i,j}^{(\ell)} - W_{0,i,j}^{(\ell)}|$ under the significance level $\alpha/2$, $\ell = 1, 2$.

(5b) Compute the critical value $\widehat{c}^{(\ell)}(q, d+1)$ of $\max_{(i,j) \in \mathcal{S}(q,d+1,\widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{i,j}^{(\ell)} - W_{0,i,j}^{(\ell)}|$ under the significance level $\alpha/2$, $\ell = 1, 2$.

Output: Decision.

(6a) Reject $H_0(0, q)$ if $\widehat{\mathbf{B}}_{q,0}^{*(\ell)} \{n^{-1/2} \widehat{c}(0, q)\} = 1$. Denote this decision by $\mathcal{D}^{(\ell)}(0, q)$.

(6b) Reject $H_0(q, d+1)$ if $\widehat{\mathbf{B}}_{d+1,q}^{*(\ell)} \{n^{-1/2} \widehat{c}(q, d+1)\} = 1$. Denote this decision by $\mathcal{D}^{(\ell)}(q, d+1)$.

(6c) Reject $H_0(q)$ if both $\mathcal{D}^{(\ell)}(0, q)$ and $\mathcal{D}^{(\ell)}(q, d+1)$ reject, for at least one $\ell = 1, 2$.

When $\widehat{W}_{j_{t+1}, j_t}^{(\ell)}$ is nonzero, the mediators j_t and j_{t+1} satisfy that $j_t \in \text{ACT}(q_2, \widehat{\mathbf{B}}^{(\ell)})$, $j_{t+1} \in \text{ACT}(q_2, \widehat{\mathbf{B}}^{(\ell)}) \cup \{q_2\}$, $q_1 \in \text{ACT}(j_t, \widehat{\mathbf{B}}^{(\ell)}) \cup \{j_t\}$, $q_1 \in \text{ACT}(j_{t+1}, \widehat{\mathbf{B}}^{(\ell)})$ and $\widehat{B}_{j_{t+1}, j_t}^{(\ell)} = 1$. Then,

$$\sqrt{|\mathcal{I}_\ell^c|} (\widehat{\mathbf{W}}^{*(\ell)})_{q_1, q_2} \leq \max_{(i,j) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{i,j}^{(\ell)} - W_{0,i,j}^{(\ell)}|, \quad (10)$$

where $\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)}) = \{(i, j) : j \in \text{ACT}(q_2, \widehat{\mathbf{B}}^{(\ell)}), i \in \text{ACT}(q_2, \widehat{\mathbf{B}}^{(\ell)}) \cup \{q_2\}, q_1 \in \text{ACT}(j, \widehat{\mathbf{B}}^{(\ell)}), \text{ or } j = q_1, q_1 \in \text{ACT}(i, \widehat{\mathbf{B}}^{(\ell)}), \{\widehat{\mathbf{B}}^{(\ell)}\}_{i,j} \neq 0\}$. Here $\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})$ denotes the set of indices

such that $\{\widehat{\mathbf{W}}^{*(\ell)}\}_{q_1, q_2}$ depends on $\widehat{\mathbf{W}}^{(\ell)}$ only through its entries in $\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})$. Then, based on (10), we use bootstrap to obtain the critical values of

$$\max_{(j_1, j_2) \in \mathcal{S}(0, q, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2}^{(\ell)}| \quad \text{and} \quad \max_{(j_1, j_2) \in \mathcal{S}(q, d+1, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2}^{(\ell)}|,$$

under the significance level $\alpha/2$. Denote the two critical values by $\widehat{c}^{(\ell)}(0, q)$ and $\widehat{c}^{(\ell)}(q, d+1)$, respectively. See Section 4.4 for more details.

Once obtaining the critical values, we reject $H_0(0, q)$ if $\widehat{\mathbf{B}}_{q,0}^{*(\ell)} \{|\mathcal{I}_\ell^c|^{-1/2} \widehat{c}^{(\ell)}(0, q)\} = 1$, and reject $H_0(q, d+1)$ if $\widehat{\mathbf{B}}_{d+1,q}^{*(\ell)} \{|\mathcal{I}_\ell^c|^{-1/2} \widehat{c}^{(\ell)}(q, d+1)\} = 1$. We reject the null $H_0(q)$ when $H_0(0, q)$ and $H_0(q, d+1)$ are both rejected. Note that, for each half of the data $\ell = 1, 2$, we have made a decision $\mathcal{D}^{(\ell)}$ regarding $H_0(q)$. Finally, we reject $H_0(q)$ when either $\mathcal{D}^{(1)}$ or $\mathcal{D}^{(2)}$ decides to reject. By Bonferroni's inequality, this yields a valid α -level test.

We summarize the full testing procedure in Algorithm 1.

4.2 Initial DAG estimation

There are multiple estimation methods available to produce an initial estimator for \mathbf{W}_0 , for instance, Zheng et al. (2018) and Yuan et al. (2019). We employ the method of Zheng et al. (2018) in our implementation. Specifically, we seek $\min_{\mathbf{W} \in \mathbb{R}^{(d+2) \times (d+2)}} \mathcal{L}(\mathbf{W}) + \lambda |\mathcal{I}_\ell| \sum_{i,j} |W_{i,j}|$, subject to $\mathbb{G}(\mathbf{W}) \in \text{DAGs}$, where $\mathcal{L}(\mathbf{W}) = \sum_{i \in \mathcal{I}_\ell} \|\widetilde{\mathbf{X}}_i - \mathbf{W} \widetilde{\mathbf{X}}_i\|_2^2$, $\lambda > 0$ is a regularization parameter, \mathbb{G} denotes the graph induced by \mathbf{W} , $\widetilde{\mathbf{X}}_i = \mathbf{X}_i - \widehat{\boldsymbol{\mu}}$ is the centered covariate, $\widehat{\boldsymbol{\mu}} = \sum_{i=1}^n \mathbf{X}_i / n$, and $|\mathcal{I}_\ell|$ is the number of data samples in the data split \mathcal{I}_ℓ . This optimization problem is challenging to solve due to the fact that the search space of DAGs scales super-exponentially with the dimension d . To resolve this issue, Zheng et al. (2018) proposed a novel characterization of the acyclic constraint, by showing that the DAG constraint can be represented by $\text{trace}\{\exp(\mathbf{W} \circ \mathbf{W})\} = d + 2$, where \circ denotes the Hadamard product, $\exp(\mathbf{A})$ is the matrix exponential of \mathbf{A} , and $\text{trace}(\mathbf{A})$ is the trace of \mathbf{A} . Then the problem becomes

$$\min_{\mathbf{W} \in \mathbb{R}^{(d+2) \times (d+2)}} \mathcal{L}(\mathbf{W}) + \lambda |\mathcal{I}_\ell| \sum_{i,j} |W_{i,j}|, \quad \text{subject to } \text{trace}\{\exp(\mathbf{W} \circ \mathbf{W})\} = d + 2. \quad (11)$$

Let $\widetilde{\mathbf{W}}^{(\ell)}$ denote the minimizer of (11). Zheng et al. (2018) proposed an efficient augmented Lagrangian based algorithm to solve (11). After obtaining $\widetilde{\mathbf{W}}^{(\ell)}$, we set the elements in its first row and last column to zero, following Condition (A2).

We make some remarks. First, the optimization in (11) is nonconvex, and there is no guarantee that the algorithm of Zheng et al. (2018) can find the global minimizer. As such, $\widetilde{\mathbf{W}}^{(\ell)}$ may not satisfy the acyclicity condition, although a global minimizer does. To meet the acyclicity constraint, Zheng et al. (2018) employed an additional thresholding step to truncate all the elements in the numerical solution to (11) whose absolute values are smaller than some threshold value c_0 to zero. We follow their implementation and adopt the same thresholding value $c_0 = 10^{-3}$. Second, to achieve the theoretical guarantees of our proposed test, we only require the estimator $\widetilde{\mathbf{W}}^{(\ell)}$ to be a consistent estimator of \mathbf{W}_0 , which is much weaker than the requirement of the test of Chakraborty et al. (2018) that the DAG estimator has to be selection consistent. In Section 5.2, we show that the global solution of (11) satisfies this consistency requirement (see Proposition 1). Meanwhile, as long as $c_0 \ll \sqrt{n^{-1} \log n}$, we can show Proposition 1 holds for the thresholded solution of (11) as well.

4.3 Screening and debiasing

Given the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$, we next compute the binary estimator $\widehat{\mathbf{B}}^{(\ell)}$ for \mathbf{B}_0 using (7) with $c = 0$. We then use the nonzero entries of $\widehat{\mathbf{B}}^{(\ell)}$ to determine the support of the decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ in the subsequent step of decorrelated estimation. As such, it serves as a screening step, and allows us to reduce the number of potential paths to a moderate level. As shown in (10), the decorrelated estimator $(\widehat{\mathbf{W}}^{*(\ell)})_{q_1, q_2}$ depends on $\widehat{\mathbf{W}}^{(\ell)}$ only through its entries in $\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})$. Consequently, the screening through $\widehat{\mathbf{B}}^{(\ell)}$ reduces the variance of $(\widehat{\mathbf{W}}^{*(\ell)})_{q_1, q_2}$, which in turn leads to an increased power for our test.

Next, we employ the decorrelated estimation idea of Ning and Liu (2017) to compute a decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ to reduce the bias of the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ obtained from (11). Because of the presence of the regularization term in (11) for high-dimensional mediators, the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ may suffer from a large bias and does not have a tractable limiting distribution. To address this issue, we refit W_{0, j_1, j_2} for any (j_1, j_2) such that $\widehat{B}_{j_1, j_2}^{(\ell)} \neq 0$, by constructing an estimating equation based on a decorrelated score function. This effectively alleviates the bias, and the resulting decorrelated estimator $\widehat{W}_{j_1, j_2}^{(\ell)}$ is both \sqrt{n} -consistent and asymptotically normal.

More specifically, after some calculations, we have that,

$$\begin{aligned} & \mathbb{E} \left(X_{j_1} - \sum_{j \neq j_2} W_{0,j_1,j} X_j \right) \left(X_{j_2} - \mathbb{E} [X_{j_2} | \{X_j\}_{j \in \text{ACT}(j_2, \mathbf{w}_0)}] \right) \\ &= W_{0,j_1,j_2} \mathbb{E} \left\{ X_{j_2} \left(X_{j_2} - \mathbb{E} [X_{j_2} | \{X_j\}_{j \in \text{ACT}(j_2, \mathbf{w}_0)}] \right) \right\}, \end{aligned} \quad (12)$$

which is the estimating equation to construct our decorrelated estimator \widehat{W}_{0,j_1,j_2} . Toward that end, we need to estimate $\mathbb{E} [X_{j_2} | \{X_j\}_{j \in \text{ACT}(j_2, \mathbf{w}_0)}]$ and $\{W_{0,j_1,j} : j \neq j_2\}$.

To estimate $\mathbb{E} [X_{j_2} | \{X_j\}_{j \in \text{ACT}(j_2, \mathbf{w}_0)}]$, we first estimate the set of ancestors of the j_2 th node $\text{ACT}(j_2, \mathbf{w}_0)$ by $\text{ACT}(j_2, \widetilde{\mathbf{w}}^{(\ell)}) = \{1 \leq j \leq d : (\widetilde{\mathbf{w}}^{*(\ell)})_{j_2,j} \neq 0\}$, for $j_2 = 1, \dots, d+1$, where $\widetilde{\mathbf{w}}^{*(\ell)} = |\widetilde{\mathbf{w}}^{(\ell)}| \oplus |\widetilde{\mathbf{w}}^{(\ell)}|^{(2)} \oplus \dots \oplus |\widetilde{\mathbf{w}}^{(\ell)}|^{(d)}$. We also note that, when estimating the ancestors, we always include the exposure variable $E = X_0$ in the set of ancestors, and always include all mediators when estimating the ancestors of the outcome variable $Y = X_{d+1}$. Next, we approximate $\mathbb{E} [X_{j_2} | \{X_j\}_{j \in \text{ACT}(j_2, \widetilde{\mathbf{w}}^{(\ell)})}]$ using a linear regression model, where the regression coefficients are estimated by,

$$\widehat{\beta}_{j_1,j_2}^{(\ell)} = \arg \min_{\substack{\beta: \beta_{j_2}=0 \\ \text{supp}(\beta) \in \text{ACT}(j_1, \widetilde{\mathbf{w}}^{(\ell)})}} \left\{ \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left(\tilde{X}_{i,j_2} - \beta^\top \widetilde{\mathbf{X}}_i \right)^2 + \sum_{\substack{k: k \neq j_2, \\ k \in \text{ACT}(j_1, \widetilde{\mathbf{w}}^{(\ell)})}} p_\lambda(|\beta_k|) \right\}, \quad (13)$$

where $\text{supp}(\beta)$ denotes the support of $\beta \in \mathbb{R}^{d+2}$, and the regression fitting is done based on the complement set of samples \mathcal{I}_ℓ^c . We choose the MCP (Zhang, 2010) penalty function and tune the penalty parameter by the Bayesian information criterion in our implementation. Alternatively, we can use LASSO (Tibshirani, 1996), SCAD (Fan and Li, 2001) or Dantzig selector (Candes and Tao, 2007) in (13). It is crucial to note that, the resulting estimator is \sqrt{n} -consistent regardless of whether the linear model approximation holds or not.

To estimate $\{W_{0,j_1,j} : j \neq j_2\}$, we employ a refined version of the initial estimator $\widetilde{\mathbf{w}}^{(\ell)}$. That is, we update the j th row of $\widetilde{\mathbf{w}}^{(\ell)}$, $j = 1, \dots, d$, by fitting a penalized regression with $\{X_{i,j}\}_{i \in \mathcal{I}_\ell}$ being the response and $\{X_{i,k}\}_{i \in \mathcal{I}_\ell, \widetilde{w}_{j,k}^{(\ell)} \neq 0}$ being the predictors. We again use the MCP penalty. Denote the resulting refined estimator by $\overline{\mathbf{w}}^{(\ell)}$. The purpose of this refitting is to improve the estimation efficiency of the initial estimator $\widetilde{\mathbf{w}}^{(\ell)}$. In our numerical experiments, we find $\overline{\mathbf{w}}^{(\ell)}$ usually converges faster than $\widetilde{\mathbf{w}}^{(\ell)}$ to the truth.

Built on the above estimators and the estimating equation (12), we debias $\widetilde{W}_{j_1,j_2}^{(\ell)}$ using

the other half of the data $\{\mathbf{X}_i\}_{i \in \mathcal{I}_\ell^c}$, for any entry such that $\widehat{B}_{j_1, j_2}^{(\ell)} \neq 0$, by

$$\widehat{W}_{j_1, j_2}^{(\ell)} = \frac{\sum_{i \in \mathcal{I}_\ell^c} \left(\widetilde{X}_{i, j_2} - \widehat{\boldsymbol{\beta}}_{j_1, j_2}^{(\ell)\top} \widetilde{\mathbf{X}}_i \right) \left(\widetilde{X}_{i, j_1} - \sum_{j \neq j_2} \widetilde{X}_{i, j} \overline{W}_{j_1, j}^{(\ell)} \right)}{\sum_{i \in \mathcal{I}_\ell^c} \widetilde{X}_{i, j_2} \left(\widetilde{X}_{i, j_2} - \widehat{\boldsymbol{\beta}}_{j_1, j_2}^{(\ell)\top} \widetilde{\mathbf{X}}_i \right)}. \quad (14)$$

We remark that, we have used the cross-fitting strategy in both the estimation of $\widehat{\boldsymbol{\beta}}_{j_1, j_2}^{(\ell)}$ in (13), and in the decorrelated estimation of $\widehat{W}_{j_1, j_2}^{(\ell)}$ in (14). This strategy guarantees each entry of the decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ is asymptotically normal, regardless of whether the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ is selection consistent or not.

4.4 Bootstrap for critical values

We next develop a multiplier bootstrap method to obtain the critical values, and summarize this procedure in Algorithm 2. Our goal is to approximate the limiting distribution of $\widehat{S}^{(\ell)} = \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2}|$ on the right-hand-side of (10).

We first observe that $\sqrt{|\mathcal{I}_\ell^c|} (\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2})$ is asymptotically equivalent to

$$\eta_{j_1, j_2}^{(\ell)} = \frac{\sqrt{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \widetilde{X}_{i, j_2} - \widehat{\boldsymbol{\beta}}^{(\ell)\top}(j_1, j_2) \widetilde{\mathbf{X}}_i \right\} \varepsilon_{i, j_1}}{\sum_{i \in \mathcal{I}_\ell^c} \widetilde{X}_{i, j_2} \left\{ \widetilde{X}_{i, j_2} - \widehat{\boldsymbol{\beta}}^{(\ell)\top}(j_1, j_2) \widetilde{\mathbf{X}}_i \right\}}. \quad (15)$$

Correspondingly, $\widehat{S}^{(\ell)} = \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell)}| + o_p(1)$, for any $q_1 \in \{0, 1, \dots, d\}$ and $q_2 \in \{1, \dots, d+1\}$. Conditioning on $\left\{ X_{i, j} : i \in \mathcal{I}_\ell^c, j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)}) \right\}$, $\eta_{j_1, j_2}^{(\ell)}$ corresponds to a sum of independent mean zero random variables, and is asymptotically normal. A rigorous proof is given in Step 1 of the proof of Theorem 1 in the appendix. This implies that $\sqrt{|\mathcal{I}_\ell^c|} (\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2})$ is asymptotically normal. Therefore, $\widehat{S}^{(\ell)}$ is to converge in distribution to a maximum of normal random variables in absolute values. Its quantile can be consistently estimated by a multiplier bootstrap method (Chernozhukov et al., 2013a).

More specifically, one can generate the bootstrap samples by replacing the residual term $\sigma_*^{-1}(\widetilde{X}_{i, j_1} - \sum_j \widetilde{X}_{i, j} \widetilde{W}_{i, j}^{(\ell)})$ in (14) with i.i.d. standard normal noise $\{e_{i, j} : 1 \leq i \leq n, 0 \leq j \leq$

Algorithm 2 Bootstrap procedure to obtain the critical values.

Input: The data $\{\mathbf{X}_i : i \in \mathcal{I}_\ell^{(c)}\}$, the significance level α , the variance estimator $\hat{\sigma}_*^2$, the number of bootstrap samples m , the estimator $\hat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2)$ from (13), and the set $\mathcal{S}(q_1, q_2, \hat{\mathbf{B}}^{(\ell)})$.

Step 1. Generate i.i.d. standard normal random variables $\{e_{i,j}^{(b)}\}_{i,j}$, $b = 1, 2, \dots, m$.

Step 2. Compute $\eta_{j_1, j_2}^{(\ell, b)*}$ according to (16), with e_{i, j_1} replaced by $e_{i, j_1}^{(b)}$, and

$$T^{(\ell, b)}(q_1, q_2) = \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \hat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell, b)*}|, \quad b = 1, 2, \dots, m.$$

Output: The empirical upper α th quantile of $\{T^{(\ell, b)}(q_1, q_2) : b = 1, \dots, m\}$.

$d + 1\}$ that are independent of the data. That is, we approximate $\eta_{j_1, j_2}^{(\ell)}$ in (15) by

$$\eta_{j_1, j_2}^{*(\ell)} = \frac{\sqrt{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \tilde{X}_{i, j_2} - \hat{\boldsymbol{\beta}}^{(\ell)\top}(j_1, j_2) \tilde{\mathbf{X}}_i \right\} e_{i, j_1} \hat{\sigma}_*}{\sum_{i \in \mathcal{I}_\ell^c} \tilde{X}_{i, j_2} \left\{ \tilde{X}_{i, j_2} - \hat{\boldsymbol{\beta}}^{(\ell)\top}(j_1, j_2) \tilde{\mathbf{X}}_i \right\}}, \quad (16)$$

where $\hat{\sigma}_*$ is some consistent estimator of σ_* . We propose to estimate σ_*^2 by $\hat{\sigma}_*^2 = \{n(d + 2)\}^{-1} \sum_{\ell \in \{1, 2\}} \sum_{i \in \mathcal{I}_\ell^c} \sum_{j=0}^{d+1} |\tilde{X}_{i, j} - \overline{\mathbf{W}}_j^{(\ell)\top} \tilde{\mathbf{X}}_i|_2^2$, where $\overline{\mathbf{W}}_j^{(\ell)}$ denotes the j th row of $\overline{\mathbf{W}}^{(\ell)}$. This estimation utilizes sample splitting again, which alleviates potential bias of the variance estimator resulting from the high correlations between the noises and the mediators in the high-dimensional setting (Fan and Lv, 2008). Lemma 4 in Section .4 of the appendix shows that $\hat{\sigma}_*^2$ is indeed consistent. Then the limiting distribution of $\hat{S}^{(\ell)}$ can be well approximated by the conditional distribution of the bootstrap samples given the data. A formal justification is given in Step 3 of the proof of Theorem 1 in the appendix.

4.5 False discovery rate control

We next present a multiple testing procedure for simultaneous inference of multivariate mediators with a proper FDR control. We present the full procedure in Algorithm 3, which consists of four steps. We next detail each step. Let \mathcal{N} be the set of unimportant mediators and \mathcal{H} be the set of our selected mediators. The FDR is defined as the expected proportion of falsely selected mediators, i.e., $\text{FDR}(\mathcal{H}) = \mathbb{E} \{ |\mathcal{N} \cup \mathcal{H}| / \max(1, |\mathcal{H}|) \}$.

Algorithm 3 Multiple testing procedure for inference of multivariate mediators.

Input: The significance level α , and the thresholding values $0 < c^{(1)}, c^{(2)} < 1$.

Step 1. Compute the p -values of testing the null hypothesis $H_0(q_1, q_2)$ for $q_1 = 0, \dots, d, q_2 = 1, \dots, d+1$ using (17) for each half of the data, $\ell = 1, 2$.

Step 2. Screening based on the pairwise minimum p -values, $\hat{p}_{\min}^{(\ell)}(q) = \min \{\hat{p}^{(\ell)}(0, q), \hat{p}^{(\ell)}(q, d+1)\}$. Let $\mathcal{H}_0^{(\ell)} = \{1 \leq q \leq d : \hat{p}_{\min}^{(\ell)}(q) \leq c^{(\ell)}\}$ denote the set of the initially selected mediators, $\ell = 1, 2$.

Step 3. Order by the pairwise maximum p -values, $\hat{p}_{\max}^{(\ell)}(q) = \max \{\hat{p}^{(\ell)}(0, q), \hat{p}^{(\ell)}(q, d+1)\}$, for those mediators in $\mathcal{H}_0^{(\ell)}$, as $\hat{p}_{(1)}^{(\ell)} \leq \hat{p}_{(2)}^{(\ell)} \leq \dots \leq \hat{p}_{(|\mathcal{H}_0^{(\ell)}|)}^{(\ell)}$.

Step 4. Select $h^{(\ell)}$ mediators in $\mathcal{H}_0^{(\ell)}$ with the smallest p -values. Let $\mathcal{H}^{(\ell)}$ denote the set of selected mediators, $\ell = 1, 2$.

Output: $\mathcal{H} = \mathcal{H}^{(1)} \cup \mathcal{H}^{(2)}$.

First, we compute the p -value of testing $H_0(q_1, q_2)$, $q_1 = 0, \dots, d, q_2 = 1, \dots, d+1$, for each half of the data. Specifically, we compute the decorrelated estimator $\widehat{\mathbf{W}}^{(\ell)}$ in Step 4 of Algorithm 1, and $\widehat{\mathbf{W}}^{*(\ell)} = |\widehat{\mathbf{W}}^{(\ell)}| \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(2)} \oplus \dots \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(d)}$. We next compute $T^{(\ell, b)}(q_1, q_2)$, $b = 1, \dots, m$, in Step 2 of Algorithm 2. Then the p -value of testing $H_0(q_1, q_2)$ in (5) is

$$\hat{p}^{(\ell)}(q_1, q_2) = \frac{1}{m} \sum_{b=1}^m \mathbb{I} \left\{ T^{(\ell, b)}(q_1, q_2) \geq \sqrt{\mathcal{I}_\ell^c}(\widehat{\mathbf{W}}^{*(\ell)})_{q_2, q_1} \right\}. \quad (17)$$

The p -values of testing $H_0(0, q)$ and $H_0(q, d+1)$ are $\hat{p}^{(\ell)}(0, q)$ and $\hat{p}^{(\ell)}(q, d+1)$, respectively.

Next, we adopt and extend the ScreenMin procedure proposed by Djordjilović et al. (2019) to our setting. We begin by computing the pairwise minimum p -values, $\hat{p}_{\min}^{(\ell)}(q) = \min \{\hat{p}^{(\ell)}(0, q), \hat{p}^{(\ell)}(q, d+1)\}$. We then screen and select those mediators whose corresponding $\hat{p}_{\min}^{(\ell)}(q)$ is smaller than a thresholding value $c^{(\ell)}$, which is determined adaptively by $c^{(\ell)} = \max \{c \in (\alpha/d, \dots, \alpha/2, \alpha) : c|\mathcal{H}_0^{(\ell)}(c)| \leq \alpha\}$, and $\mathcal{H}_0^{(\ell)}(c)$ denotes the set of pre-screened mediators when the threshold value is c . Djordjilović et al. (2019) showed such a thresholding value approximately maximizes the power to reject false union hypotheses. It also works well in our numerical studies. Denote the resulting set of important mediators by the ScreenMin procedure as $\mathcal{H}_0^{(\ell)}$.

Next, we compute the pairwise maximum p -value, which is also the p -value of testing the significance of an individual mediator $H_0(q)$ in our setting, $\hat{p}_{\max}^{(\ell)}(q) = \max \{\hat{p}^{(\ell)}(0, q), \hat{p}^{(\ell)}(q, d+1)\}$. We order the mediators in $\mathcal{H}_0^{(\ell)}$ according to $\hat{p}_{\max}^{(\ell)}(q)$.

Finally, we apply the procedure of Benjamini and Yekutieli (2001a) to the ordered mediators, and select $h^{(\ell)}$ mediators with the smallest p -values, where $h^{(\ell)} = \max \left[i : \widehat{p}_{(i)}^{(\ell)} \leq (i\alpha) / \{2|\mathcal{H}_0^{(\ell)}| \sum_{j=1}^{|\mathcal{H}_0^{(\ell)}|} j^{-1}\} \right]$. Letting $\mathcal{H}^{(\ell)}$ denote the selected mediators for each half of the data, $\ell = 1, 2$, respectively, we set the final set of selected mediators as $\mathcal{H} = \mathcal{H}^{(1)} \cup \mathcal{H}^{(2)}$.

5 Theory

In this section, we first establish the consistency of our test for each individual mediator, by deriving the asymptotic size and power. We then show that the multiple testing procedure achieves a valid FDR control. Finally, as a by-product, we derive an oracle inequality for the estimator $\widetilde{\mathbf{W}}^{(\ell)}$ computed from (11) using the method of Zheng et al. (2018).

5.1 Consistency and FDR control

We first present a main regularity condition (A4), while we defer two additional regularity conditions (A5) and (A6) to Section .3 of the appendix in the interest of space.

(A4) With probability approaching one, $\text{ACT}(j, \widetilde{\mathbf{W}}^{(\ell)})$ contains all parents of j , for any $j = 0, \dots, d+1, \ell = 1, 2$.

This condition requires an appropriate identification of the graph, in that $\text{ACT}(j, \widetilde{\mathbf{W}}^{(\ell)})$ contains all parents of j . It serves as a basis for the asymptotic properties of the proposed test. We make some remarks. First, this condition is weaker than requiring $\widetilde{\mathbf{W}}_j^{(\ell)}$, the j th column of $\widetilde{\mathbf{W}}^{(\ell)}$, to satisfy the sure screening property; i.e, $\text{PA}(j) \subseteq \text{supp}(\widetilde{\mathbf{W}}_j^{(\ell)})$, where $\text{PA}(j)$ denotes the parents of node j . To better illustrate this, consider the DAG example in Figure 1(b). For node $j = 3$, the sure screening property of $\widetilde{\mathbf{W}}_3^{(\ell)}$ requires $W_{0,3,1}$ and $W_{0,3,2}$ to satisfy certain minimum-signal-strength conditions; see, e.g., Fan and Lv (2008). In comparison, we require $\{1, 2\} \in \text{ACT}(3, \widetilde{\mathbf{W}}^{(\ell)})$ with probability approaching one. This requires either $\{W_{0,3,1}, W_{0,3,2}\}$, or $\{W_{0,2,1}, W_{0,3,2}\}$, to satisfy certain minimum-signal-strength conditions. In that sense, our test is “doubly robust”. In Proposition 1, we show (A4) is satisfied when the initial estimator is obtained using (11) of Zheng et al. (2018). Second, we remark that, when (A4) is not satisfied, the bias in the ancestor identification step is to affect the subsequent testing procedure. This phenomenon is similar to post-selection inference in linear regressions, where direct inference may fail when the variable selection step does

not satisfy the sure screening or selection consistency property (Meinshausen et al., 2009; Shi et al., 2019; Zhu et al., 2020). In the linear regression setting, debiasing is an effective remedy to address the issue. However, the usual debiasing strategy may not be directly applicable in our setting; see Section .3 of the appendix for more discussion. We leave this post-selection inference problem as future research.

We next establish the validity of our test for a single mediator in Theorem 1, and its local power property in Theorem 2. Combining the two theorems yields its consistency.

Theorem 1. *Suppose (A1) to (A5) hold. Suppose $d = O(n^{\kappa_1})$ for some constant $\kappa_1 > 0$, and $\|\mathbf{W}_0\|_2$ is bounded. Then for a significance level $0 < \alpha < 1$, and any mediator $q = 1, \dots, d$, the proposed test in Algorithm 1 satisfies that*

$$\Pr\left\{H_0(q) \text{ is rejected} \mid H_0(q) \text{ holds}\right\} \leq \alpha + o(1).$$

Next, for any directed path $\zeta: E \rightarrow M_{i_1} \rightarrow \dots \rightarrow M_{i_k} \rightarrow Y$, define ω_ζ^* as the minimum signal strength along this path,

$$\omega_\zeta^* = \min \left\{ |W_{0,i_1,0}|, \min_{j \in \{1, \dots, k-1\}} |W_{0,i_{j+1},i_j}|, |W_{0,d+1,i_k}| \right\}.$$

Under the alternative hypothesis $H_1(q)$, there exists at least one path ζ that passes through M_q such that $\omega_\zeta^* > 0$. We next establish the local power property of our test.

Theorem 2. *Suppose the conditions in Theorem 1 hold. Suppose $\max_{j \in \{0,1,\dots,d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 = O_p(n^{-1/2}\sqrt{\log n})$, where $\widetilde{\mathbf{W}}_j^{(\ell)}$ is the j th row of $\widetilde{\mathbf{W}}_j$. Suppose there exists one path $\zeta: E \rightarrow M_{i_1} \rightarrow \dots \rightarrow M_{i_k} \rightarrow Y$ that passes through M_q such that $\omega_\zeta^* \gg n^{-1/2}\sqrt{\log n}$ under $H_1(q)$. Then the proposed test in Algorithm 1 satisfies that,*

$$\Pr\left\{H_0(q) \text{ is rejected} \mid H_1(q) \text{ holds}\right\} \rightarrow 1, \quad \text{as } n \rightarrow \infty.$$

Note that we require $\omega_\zeta^* \gg n^{-1/2}\sqrt{\log n}$ for some ζ in Theorem 2. Consequently, our test is consistent against some local alternatives that are \sqrt{n} -consistent to the null up to some logarithmic term. Let $s_0 = \max_j |\text{supp}(\mathbf{W}_{0,j})|$ denotes the maximum sparsity size where $\mathbf{W}_{0,j}$ stands for the j th row of \mathbf{W}_0 . In Proposition 1, we show that $\max_{j \in \{0,\dots,d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_j\|_2 = O_p(n^{-1/2}\sqrt{\log n})$ when the maximum sparsity size s_0 is bounded.

Next, we show that our multiple testing procedure achieves a valid FDR control. Note that we use a union-intersection principle to construct the p -value for $H_0(q)$. The key idea

of the ScreenMin procedure of Djordjilović et al. (2019) lies in exploiting the independence between the two p -values $\hat{p}^{(\ell)}(0, q)$ and $\hat{p}^{(\ell)}(q, d + 1)$. In our setting, these p -values are actually asymptotically independent. We thus have the following result.

Theorem 3. *Suppose the conditions in Theorem 1 hold. Then the set of selected mediators \mathcal{H} in Algorithm 3 satisfies that $FDR(\mathcal{H}) \leq \alpha + o(1)$.*

5.2 Oracle inequality for the initial DAG estimator

As a by-product, we establish the oracle inequality for the estimator of Zheng et al. (2018). We first introduce the oracle estimator. For a given ordering $\pi = \{\pi_0, \pi_1, \dots, \pi_{d+1}\}$, consider the estimator $\widetilde{\mathbf{W}}^{(\ell)}(\pi) = \left\{ \widetilde{\mathbf{W}}_0^{(\ell)}(\pi), \widetilde{\mathbf{W}}_1^{(\ell)}(\pi), \dots, \widetilde{\mathbf{W}}_{d+1}^{(\ell)}(\pi) \right\}^\top$ where

$$\widetilde{\mathbf{W}}_{\pi_j}^{(\ell)}(\pi) = \arg \min_{\beta: \text{supp}(\beta) \in \{\pi_0, \pi_1, \dots, \pi_{j-1}\}} \sum_{i \in \mathcal{I}_\ell} (X_{i, \pi_j} - \beta^\top \mathbf{X}_i)^2 + \lambda |\mathcal{I}_\ell| \|\beta\|_1,$$

for $j \in \{0, 1, \dots, d + 1\}$. That is, $\widetilde{\mathbf{W}}^{(\ell)}(\pi)$ is computed as if the ordering of π were known. Let Π^* denote the set of all true orderings, while a more rigorous definition is given in Section .3 of the supplementary appendix. Then the oracle estimator $\widetilde{\mathbf{W}}_{\pi_j}^{(\ell)}(\pi^*)$, for some $\pi^* \in \Pi^*$, is computed as if the true ordering π^* were known. With a proper choice of λ , it follows from the oracle inequality for LASSO (Bickel et al., 2009) that,

$$\max_{j \in \{0, \dots, d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi^*) - \mathbf{W}_{0,j}\|_2 \leq O(1)n^{-1/2}\sqrt{s_0 \log n}.$$

The next proposition establishes the convergence rate of $\widetilde{\mathbf{W}}^{(\ell)}$ obtained from (11).

Proposition 1. *Suppose (A1), (A2), (A3) and (A6) hold. Suppose $d = O(n^{\kappa_1})$ for some constant $\kappa_1 < 1$, $\|\mathbf{W}_0\|_2$ is bounded, and $\lambda = \kappa_2 n^{-1/2} \sqrt{\log n}$ for some sufficiently large constant $\kappa_2 > 0$. Then with probability tending to 1, the initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ obtained from (11) satisfies that,*

$$\widetilde{\mathbf{W}}^{(\ell)} = \widetilde{\mathbf{W}}^{(\ell)}(\pi^*) \text{ for some } \pi^* \in \Pi^*, \text{ and } \max_{j \in \{0, \dots, d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 \leq O(1)n^{-1/2}\sqrt{s_0 \log n}.$$

Proposition 1 shows that the convergence rate of $\widetilde{\mathbf{W}}^{(\ell)}$ is the same as that of the oracle estimator. Moreover, the true ordering π^* can be inferred from $\widetilde{\mathbf{W}}^{(\ell)}$. If we produce the initial estimator from (11), it further implies that (A4) holds. We again make some remarks.

First, in this proposition, we require the dimension d to grow at a slower rate than n . We note that this condition is not needed in Theorems 1-3. Moreover, we can further relax this requirement on d by imposing some sparsity conditions on $\widetilde{\mathbf{W}}^{(\ell)}$ and the population limit of $\widetilde{\mathbf{W}}^{(\ell)}(\pi)$; see the final remark of the proof of Proposition 1 in the appendix. Second, this proposition is for the global minimizer of (11). Of course, as we have commented, there is no guarantee that the algorithm of Zheng et al. (2018) can find the global minimizer. This is a universal problem for almost all nonconvex optimizations. Nevertheless, Zhong et al. (2014, Theorem 1) showed that the actually minimizer to (11), which is obtained through the proximal quasi-Newton method, converges to the local minimizer of the augmented Lagrange problem, while Zheng et al. (2018, Table 1) showed that numerically the difference between the actual minimizer and the global minimizer is much smaller than that between the global minimizer and the ground truth. As such, we expect Proposition 1 to hold for the actually minimizer as well. Meanwhile, we acknowledge that this local minimizer problem is challenging and is warranted for future research.

6 Simulations

We simulate the data following model (1). We set $\boldsymbol{\mu}_0$ to a vector of ones, and $\sigma_*^2 = 1$. We generate the adjacency matrix \mathbf{W}_0 as follows: We begin with a zero matrix, then replace every entry W_{0,j_1,j_2} in the lower off-diagonals by the product of two random variables $R_{j_1,j_2}^{(1)} R_{j_1,j_2}^{(2)}$. Here $R_{j_1,j_2}^{(1)} \sim \text{Bernoulli}(p_1)$, if $j_2 = 0$, or $j_1 = d + 1$, and $R_{j_1,j_2}^{(1)} \sim \text{Bernoulli}(p_2)$, otherwise, and $R_{j_1,j_2}^{(2)}$ is uniformly distributed on $[-2, -0.5] \cup [0.5, 2]$. All these variables are independently generated. We consider three scenarios of the total number of mediators d , with varying binary probabilities p_1, p_2 , each under two sample sizes n ; i.e., $(d, p_1, p_2) = (50, 0.05, 0.15)$ with $n = 100, 200$, $(d, p_1, p_2) = (100, 0.03, 0.1)$ with $n = 250, 500$, and $(d, p_1, p_2) = (150, 0.02, 0.05)$ with $n = 250, 500$. Table 2 in Section .10 of the appendix reports the corresponding mediators with nonzero mediation effects, and their associated $\delta(q)$, where $\delta(q) = (\mathbf{W}_0^*)_{d+1,q}(\mathbf{W}_0^*)_{q,0}$, and \mathbf{W}_0^* is constructed based on \mathbf{W}_0 , $q = 1, \dots, d$. By Lemma 2, $\delta(q)$ measures the size of the mediation effect. When $\delta(q) = 0$, $H_0(q)$ holds; otherwise, $H_1(q)$ holds. A larger $\delta(q)$ indicates a stronger mediation effect. The percentage of nonzero mediators for the three scenarios is 0.12, 0.09 and 0.06, respectively.

We first evaluate the empirical performance of our test for a single mediator in Algorithm

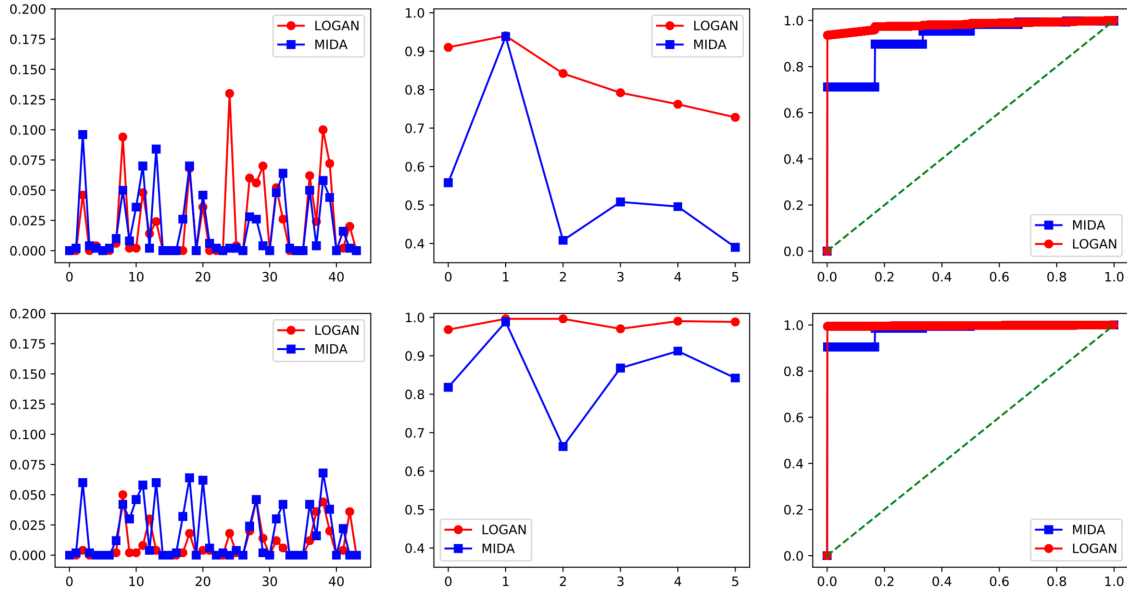


Figure 2: Empirical rejection rate and ROC curve of the proposed test, LOGAN, and the test of Chakraborty et al. (2018), MIDA, when $d = 50$. The upper panels: $n = 100$, and the bottom panels: $n = 200$. The left panels: under H_0 , the middle panels: under H_1 , where the horizontal axis is the mediator index, and the right panels: the average ROC curve.

1. We also compare it with that of Chakraborty et al. (2018), which they named as Mediation Interventional calculus when the DAG is Absent (MIDA). We use the same initial estimator as ours for MIDA. We construct the $100(1 - \alpha)\%$ confidence interval for the total effect of each mediator, following the procedure as described in Chakraborty et al. (2018). We reject the null hypothesis if zero is not covered by the confidence interval.

We evaluate each testing method by the empirical rejection rate, in percentage, out of 500 data replications at the significance level $\alpha = 5\%$. This rate reflects the size of the test when the null hypothesis holds, and reflects the power otherwise. We also compute the average receiver operating characteristic (ROC) curves, aggregated over 500 replications, when the significance level α varies. Figure 2 reports the results when $d = 50$. The results for $d = 100$ and $d = 150$ show a similar qualitative pattern, and are reported in Section .10 of the appendix. We make a few observations. First, our test achieves a valid size under the null hypothesis. The empirical rejection rate is close to or below the nominal level for most cases. When the sample size n is small, our test has a few inflated type-I errors. As n increases, all the rejection rates are below the nominal level. By contrast, the test of Chakraborty et al. (2018) still has a good number of inflated type-I errors even when n is large. Such




Figure 3: False discover rate and true positive rate of the proposed method and the Benjamini-Yekutieli procedure when $d = 50$. The horizontal axis corresponds to the significance level α . The left two panels: $n = 100$, and the right two panels: $n = 200$.

inflated errors may be due to the fact that MIDA relies on the selection consistency of the estimated DAG, which may not hold under the finite samples. Second, our test consistently achieves a larger empirical power over MIDA under the alternative hypothesis. This may be due to that the effects calculated by MIDA along different paths may cancel each other, leading to a decreased power. Combined with the results on the empirical size, the power of our test is not gained at the cost of the inflated Type-I errors. Moreover, the empirical power of our test increases along with the sample size, demonstrating the consistency of the test. Finally, we observe that the ROC curve of our test lies above that of MIDA in all settings as α varies, which clearly demonstrates the advantage of our test over MIDA.

We next evaluate the empirical performance of our multiple testing procedure in Algorithm 3. We also compare it with the standard Benjamini-Yekutieli (BY) procedure. For the latter, in Step 2 of Algorithm 3, instead of applying ScreenMin to determine the set $\mathcal{H}_0^{(\ell)}$, one simply sets $\mathcal{H}_0^{(\ell)} = \{1, 2, \dots, d\}$, i.e., the set of all mediators. We evaluate each testing procedure by the false discovery rate (FDR) and the true positive rate (TPR), over 500 data replications. Figure 3 reports the results under the varying significance level α from 0 to 0.4 when $d = 50$. The results for $d = 100$ and $d = 150$ are similar, and are reported in Section .10 of the appendix. It is seen that both methods achieve a valid false discovery control, in that the FDRs are all below the nominal level. However, our method is more powerful than BY, as reflected by a larger TPR in all cases.

7 Application

In this section, we illustrate our testing method with an application to a neuroimaging study of Alzheimer’s disease (AD). AD is an irreversible neurodegenerative disorder, and is characterized by progressive impairment of cognitive and memory functions. It is the leading form of dementia, and the sixth leading cause of death in the U.S (Alzheimer’s Association, 2020). The data we analyze is part of the ongoing Berkeley Aging Cohort Study. It consists of 698 participants aging between 55.3 and 94.1 years old. For each participant, the well established PACC composite score was recorded, which combines tests that assess episodic memory, timed executive function, and global cognition (Donohue et al., 2014). Moreover, for each participant, a 1.5T structural magnetic resonance imaging (MRI) scan and a positron emission tomography (PET) scan using 18-F florbetaben tracer were acquired. All imaging data were preprocessed following the established protocols. Particularly, for MRI, all T1 images were bias-corrected, segmented, then warped and normalized to a common template space. Then the volumes were examined quantitatively by a cortical surface-based analysis and turned into cortical thickness measures. Cortical thickness is an important biomarker that reflects AD severity. We employ the FreeSurfer brain atlas and summarize each MRI image by a 68-dimensional vector, whose entries measure cortical thickness of 68 brain regions of interest. For PET, native-space images were realigned and coregistered to each participant’s MRI scan, and centiloid analysis was performed to transform the standardized uptake value ratio to centiloid units. The PET scan provides a measure of deposition of amyloid-beta, a hallmark pathological protein of AD that is commonly found in the brains of AD and elderly subjects. The total amount of amyloid-beta deposition was extracted from PET for each subject. There are well validated methods for thresholding the subjects based on the total deposition as amyloid positive and amyloid negative groups, which are known to behave differently in AD progression (Landau et al., 2013). For our data, 309 subjects were classified as amyloid positive, and 389 as amyloid negative. Since age is a well known risk factor for AD, in our study, we aim to understand how age mediates cortical thickness of different brain regions then the PACC score. We carry out the mediation analysis for the amyloid positive and amyloid negative groups separately.

We apply the proposed multiple testing procedure in Algorithm 3 to this data, with age as the exposure, the cortical thickness of 68 brain regions as the potential mediators,

and the PACC score as the outcome. We set the FDR level at 10%. For the amyloid positive group, we find one significant mediator, and for the amyloid negative group, we found six significant mediators. Table 1 reports the results. These findings agree well with the neuroscience literature. In particular, the entorhinal cortex functions as a hub in a widespread network for memory, navigation and the perception of time. It is found implicated in the early stages of AD, and is one of the most heavily damaged cortices in AD (van Hoesen et al., 1991). The precuneus is involved with episodic memory, visuospatial processing, reflections upon self, and aspects of consciousness, and is found to be an AD-signature region (Bakkour et al., 2013). Moreover, the superior temporal gyrus is involved in auditory processing, and also has been implicated as a critical structure in social cognition. The superior frontal gyrus is involved in self-awareness, and the inferior parietal lobule is involved in the perception of emotions. Numerous studies have found involvement of these brain regions in the development of AD (Du et al., 2007; Bakkour et al., 2013).

8 Discussion

In this article, we have primarily focused on the case when there is only a single DAG associated with our model. Now, we briefly discuss the extension to the case when there is an equivalence class of DAGs. Specifically, when the error variances σ_i^2 , $i = 0, \dots, d + 1$, in (A3) are not all equal, there exist an equivalence class of DAGs, denoted by \mathcal{G} , that could generate the same joint distribution of the variables. Such a class can be uniquely represented by a completed partially directed acyclic graph. For each DAG $\mathbb{G} \in \mathcal{G}$, we define $\omega_\zeta(\mathbb{G})$ as the total effect of E on Y attributed to a given path ζ following (2). Then, our hypotheses of interest become,

$$\begin{aligned} H_0(q) : \omega_\zeta(\mathbb{G}) &= 0, \quad \text{for all } \zeta \text{ that passes through } M_q \text{ and all } \mathbb{G} \in \mathcal{G} \quad \text{versus} \\ H_1(q) : \omega_\zeta(\mathbb{G}) &\neq 0, \quad \text{for some } \zeta \text{ that passes through } M_q \text{ and some } \mathbb{G} \in \mathcal{G}. \end{aligned} \tag{18}$$

Table 1: Identified significant mediators for the amyloid positive and amyloid negative groups.

Amyloid positive group		Amyloid negative group	
r-entorhinal	l-entorhinal	l-precuneus	l-superiortemporal
	r-inferiorparietal	r-superiorfrontal	r-superiortemporal

To test (18), we begin by estimating the equivalence class \mathcal{G} based on each half of the dataset. This can be done by applying the structural learning algorithm such as Chickering (2003). Let $\widehat{\mathcal{G}}$ denote the resulting estimator. For each $\widehat{\mathbb{G}} \in \widehat{\mathcal{G}}$, we employ the procedure in Section 4 to construct a test statistic. We then take the supremum of these test statistics over all $\widehat{\mathbb{G}}$, and obtain its critical value via bootstrap.

Finally, we comment that our proposed testing procedures can be extended to more scenarios, e.g., when there are sequentially ordered multiple sets of mediators, or when there are multiple exposure variables. We can also speed up the computation of the Boolean matrices using some transition closure algorithm (Chakradhar et al., 1993) when the dimension of the DAG is large. We leave those pursuits as our future research.

Acknowledgements

Li's research was partially supported by NIH grants R01AG061303, R01AG062542, and R01AG034570. The authors thank the AE, and the reviewers for their constructive comments, which have led to a significant improvement of the earlier version of this article.

Appendix

We first outline a multi-split version of our test. We then introduce a DAG learning procedure under a weaker constant variance condition. We next present some additional regularity conditions, followed by two supporting lemmas, then the proofs of the main theoretical results in the paper. Finally, we present some additional numerical results.

We employ the following notation. For any sequence $\{a_n : n \geq 1\}$, $a_n = O(1)$ means $|a_n| \leq C$ for some constant $C > 0$, and $a_n = o(1)$ means $\lim_n a_n = 0$. For a sequence of random variables $\{Z_n : n \geq 1\}$, $Z_n = O_p(1)$ means, for any sufficiently small $\varepsilon > 0$, there exists some constant $M > 0$ such that $\Pr(|Z_n| \leq M) \geq 1 - \varepsilon$, and $Z_n = o_p(1)$ means $\{Z_n : n \geq 1\}$ converges in probability to zero. Without loss of generality, we assume $\boldsymbol{\mu}_0 = 0$. To simplify the presentation, we only consider the case where $\widehat{\boldsymbol{\mu}} = 0$ and hence $\widetilde{\mathbf{X}}_i = \mathbf{X}_i$ for $i = 1, \dots, n$. In the case where $\widehat{\boldsymbol{\mu}} \neq 0$, the theories can be similarly proved.

.1 A multi-split version of the test

We first develop a version of our individual mediator test based on multiple binary splits. This helps improve the power when the sample size is limited, and also helps mitigate the randomization arising from a single binary split. The main idea is to apply the single-split method in Algorithm 1 multiple times, then combine the p -values from all splits. Specifically, we carry out the binary split S times. For the s th binary split, we divide $\{1, \dots, n\}$ into two disjoint subsets $\mathcal{I}_{s,1} \cup \mathcal{I}_{s,2}$ of equal sizes. We then apply Algorithm 1 to compute the p -values for $H_0(0, q)$ and $H_0(q, d+1)$ for each half of the data. Denote the obtained p -values by $\hat{p}^{(s,1)}(0, q)$, $\hat{p}^{(s,1)}(q, d+1)$, $\hat{p}^{(s,2)}(0, q)$ and $\hat{p}^{(s,2)}(q, d+1)$, respectively. We next combine these p -values following the idea of Meinshausen et al. (2009), by defining

$$\begin{aligned}\hat{p}(0, q) &= \min \left(1, q_\gamma \left[\left\{ \gamma^{-1} \hat{p}^{(s,\ell)}(0, q), s = 1, \dots, S, \ell = 1, 2 \right\} \right] \right), \\ \hat{p}(q, d+1) &= \min \left(1, q_\gamma \left[\left\{ \gamma^{-1} \hat{p}^{(s,\ell)}(q, d+1), s = 1, \dots, S, \ell = 1, 2 \right\} \right] \right),\end{aligned}$$

where γ is some constant between 0 and 1, and q_γ is the empirical γ -quantile. In our simulations, we have experimented with a range of values of γ between 0.1 and 0.2, and the results are similar, and thus we set $\gamma = 0.15$ in our implementation. The corresponding p -value for $H_0(q)$ is given by $\max\{\hat{p}(0, q), \hat{p}(q, d+1)\}$, following the union-intersection principle.

We apply this multi-split method to the simulation examples in Section 6, and compare with the single-split method. Figure 4 shows the empirical rejection rates of the two methods. It is seen that the multi-split method in general improves over the single-split method, by achieving smaller type-I errors and larger powers.

.2 Learning DAG under a weaker constant variance condition

We note that the constant variance condition in (A3) can be relaxed as follows.

(A3*) The errors ε_i , $i = 0, 1, \dots, d+1$, are jointly normally distributed and independent.

In addition, the error variances $\sigma_i^2 = \text{Var}(\varepsilon_i)$, $i = 1, \dots, d$, are constant; i.e., $\sigma_1^2 = \dots = \sigma_d^2 = \sigma_*^2$ for some constant $\sigma_* > 0$.

In other words, the constant variance requirement does not have to be imposed on the exposure and outcome variables. Under this weaker requirement, we first need to ensure \mathbf{W}_0 remains identifiable, since Peters and Bühlmann (2014, Theorem 1) is no longer directly applicable. In addition, we need to modify the initial estimator of \mathbf{W}_0 accordingly.




Figure 4: Empirical rejection rate of the single-split (LOGAN-SS) and the multi-split (LOGAN-MS) method. The upper panels: $(d, n) = (50, 100)$, the middle panels: $(d, n) = (100, 250)$, and the bottom panels: $(d, n) = (150, 250)$. The left panels: under H_0 , the middle panels: under H_1 , where the horizontal axis is the mediator index, and the right panels: the average ROC curve.

The next lemma shows that \mathbf{W}_0 remains identifiable under $(A3^*)$.

Lemma 3. *Suppose $(A1)$, $(A2)$ and $(A3^*)$ hold. Then \mathbf{W}_0 is identifiable from the joint distribution function of \mathbf{X} .*

Proof: By $(A2)$, we decompose \mathbf{W}_0 as,

$$\mathbf{W}_0 = \begin{pmatrix} 0 & \mathbf{0}_d^\top & 0 \\ \mathbf{W}_{0,1} & \mathbf{W}_{1,1} & \mathbf{0}_d \\ W_{0,2} & \mathbf{W}_{1,2}^\top & 0 \end{pmatrix}, \quad (19)$$

where $W_{0,2} \in \mathbb{R}$, $\mathbf{W}_{0,1}, \mathbf{W}_{1,2} \in \mathbb{R}^d$, $\mathbf{W}_{1,1} \in \mathbb{R}^{d \times d}$, and the matrix $\mathbf{W}_{1,1}$ is acyclic under $(A1)$.

We first note that $W_{0,2}$ and $\mathbf{W}_{1,2}$ correspond to the regression coefficients of $(E, \mathbf{M})^\top$ on Y . Under the given model, the covariance matrix of (E, \mathbf{M}^\top) is non-degenerated. As

such, $W_{0,2}$ and $\mathbf{W}_{1,2}$ are uniquely determined by the distribution function of \mathbf{X} .

We next show that $\mathbf{W}_{1,1}$ is also uniquely determined by the distribution function of \mathbf{X} . For each $j = 1, \dots, d$, let \widetilde{M}_j denote the population residual adjusted by the exposure, i.e., $\widetilde{M}_j = (M_j - \mu_{0,j}) - \text{corr}(M_j, E)(E - \mu_{0,0})$. It follows that the set of residuals $\widetilde{\mathbf{M}}$ satisfy that $\widetilde{\mathbf{M}} = \mathbf{W}_{1,1}\widetilde{\mathbf{M}} + (\varepsilon_1, \dots, \varepsilon_d)^\top$. By (A1), $\mathbf{W}_{1,1}$ is acyclic. As such, $\widetilde{\mathbf{M}}$ forms a structural linear equation with the coefficient matrix $\mathbf{W}_{1,1}$. Under (A3*), all the residuals have the constant variance. It then follows from Theorem 1 of Peters and Bühlmann (2014) that $\mathbf{W}_{1,1}$ is identifiable.

Finally, we note that $\mathbf{W}_{0,1}$ satisfies $\mathbf{W}_{0,1} = \text{cov}(E, \mathbf{M} - \mathbf{W}_{1,1}\mathbf{M})/\text{Var}(E)$. It follows from the identifiability of $\mathbf{W}_{1,1}$ that $\mathbf{W}_{0,1}$ is identifiable as well.

This completes the proof. \square

We next outline a modified initial DAG estimation procedure under the new constant variance condition (A3*). Similar to Proposition 1, we can show that this new estimator satisfies the oracle inequality as well.

Following the decomposition of \mathbf{W}_0 in (19), we first estimate $\mathbf{W}_{0,2}$ and $\mathbf{W}_{1,2}$ using penalized regressions of $(E, \mathbf{M})^\top$ on Y such as MCP, LASSO, SCAD, and Dantzig selector. Denote the corresponding estimators as $\widetilde{W}_{0,2}^{(\ell)}$ and $\widetilde{\mathbf{W}}_{1,2}^{(\ell)}$.

We next estimate $\mathbf{W}_{1,1}$, by first regressing \mathbf{M}_i on E_i , $i = 1, \dots, n$, to obtain the estimated residual $\widehat{\widetilde{M}}_i$, then employing the method of Zheng et al. (2018) to solve

$$\widetilde{\mathbf{W}}_{1,1}^{(\ell)} = \arg \min_{\mathbf{W} \in \mathbb{R}^{d \times d}} \sum_{i \in \mathcal{I}_\ell} \|\widehat{\widetilde{M}}_i - \mathbf{W} \widehat{\widetilde{M}}_i\|_2^2 + \lambda |\mathcal{I}_\ell| \sum_{i,j} |W_{i,j}| \text{ subject to } \text{trace}\{\exp(\mathbf{W} \circ \mathbf{W})\} = d.$$

After obtaining $\widetilde{\mathbf{W}}_{1,1}^{(\ell)}$, we set $\widetilde{\mathbf{W}}_{0,1}^{(\ell)} = \widetilde{\mathbf{W}}_{1,1}^{(\ell)} \widehat{\text{cov}}(\mathbf{M}, E) \widehat{\text{Var}}^{-1}(E)$, where $\widehat{\text{cov}}(\mathbf{M}, E)$ is the sampling covariance estimator of \mathbf{M} , and $\widehat{\text{Var}}(E)$ is the sampling variance estimator of E .

Finally, we put together $\widetilde{\mathbf{W}}_{0,1}^{(\ell)}$, $\widetilde{\mathbf{W}}_{1,1}^{(\ell)}$, $\widetilde{W}_{0,2}^{(\ell)}$ and $\widetilde{\mathbf{W}}_{1,2}^{(\ell)}$ according to (19) to form the modified initial estimator $\widetilde{\mathbf{W}}^{(\ell)}$ for \mathbf{W}_0 .




Figure 5: An illustrative DAG with four nodes, where both $(0, 1, 2, 3)$ and $(0, 2, 1, 3)$ correspond to the true orderings.

.3 Additional regularity conditions

We introduce two additional regularity conditions for the theoretical guarantees of the proposed test. We begin with some notation. Define the limit of the estimator $\widehat{\beta}^{(\ell)}(j_1, j_2)$

$$\beta_0^{(\ell)}(j_1, j_2) = \arg \min_{\beta: \beta_{j_2}=0, \text{supp}(\beta) \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})} \mathbb{E} (X_{j_2} - \beta^\top \mathbf{X}_i)^2.$$

Any permutation $\pi = (\pi_0, \pi_1, \dots, \pi_d, \pi_{d+1})^\top$ of $\{0, 1, \dots, d, d+1\}$ determines an order of the mediators $\{M_j\}_{1 \leq j \leq d}$. Define

$$\mathbf{W}_{\pi_j}(\pi) = \arg \min_{\beta: \text{supp}(\beta) \in \{\pi_0, \pi_1, \dots, \pi_{j-1}, \pi_j\}} \mathbb{E} (M_{\pi_j} - \beta^\top \mathbf{M})^2, \quad \text{for } j = 1, \dots, d.$$

Let $\mathbf{W}(\pi) = \{\mathbf{W}_0(\pi), \mathbf{W}_1(\pi), \dots, \mathbf{W}_{d+1}(\pi)\}^\top$. It corresponds to the coefficient matrix obtained by doing a Gram-Schmidt orthogonalization, starting with X_{π_0} , and finishing by projecting $X_{\pi_{d+1}}$ on $X_{\pi_0}, X_{\pi_1}, \dots, X_{\pi_d}$. Let $\mathbf{\Omega}(\pi)$ be a diagonal matrix where the diagonal elements $\omega_0^2(\pi), \omega_1^2(\pi), \dots, \omega_{d+1}^2(\pi)$ correspond to the error variances, $\text{Var}\{X_0 - \mathbf{X}^\top \mathbf{W}_0(\pi)\}, \text{Var}\{X_1 - \mathbf{X}^\top \mathbf{W}_1(\pi)\}, \dots, \text{Var}\{X_{d+1} - \mathbf{X}^\top \mathbf{W}_{d+1}(\pi)\}$, respectively. Let Π^* denote the set consisting of all true orderings π^* such that $\mathbf{W}(\pi^*) = \mathbf{W}_0$. Note that π^* may not be unique. As an illustration, consider the DAG in Figure 5, where both $(0, 1, 2, 3)$ and $(0, 2, 1, 3)$ correspond to the true orderings of the four nodes.

We impose the following additional regularity conditions. In particular, (A5) is required to establish the consistency and FDR control of the proposed test, while (A6) is to establish the oracle inequality for the initial estimator.

(A5) There exist some constants $\kappa_3, \kappa_4, \kappa_5, \kappa_6 > 0$, with $\kappa_4 + \kappa_5 > 1/2$, such that, $\|\widehat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^{(\ell)}(j_1, j_2)\|_2 \leq \kappa_3 n^{-\kappa_4}$, $\|\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0,j_1}\|_2 \leq \kappa_3 n^{-\kappa_5}$ and $\|\widehat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^{(\ell)}(j_1, j_2)\|_1 \leq \kappa_3 n^{-\kappa_6}$, $\|\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0,j_1}\|_1 \leq \kappa_3 n^{-\kappa_6}$, for any $0 \leq j_1, j_2 \leq d+1, \ell = 1, 2$.

(A6) There exists a constant $\omega > 0$ such that for all $\pi \notin \Pi^*$, $\frac{1}{d} \sum_{j=0}^{d+1} \{\omega_j^2(\pi) - \sigma_*^2\}^2 > \omega$.

Condition (A5) is mild. This is because, when (A4) holds and $\beta_0^{(\ell)}$ is estimated via the MCP, LASSO, SCAD, or Dantzig selector, we have $\max_{j_1, j_2, \ell} \|\hat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^{(\ell)}(j_1, j_2)\|_2 \leq O(1)n^{-1/2}\sqrt{s^* \log n}$ and $\max_{j_1, j_2, \ell} \|\hat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^{(\ell)}(j_1, j_2)\|_1 \leq O(1)n^{-1/2}s^*\sqrt{\log n}$, with probability approaching one, where $s^* = \max_{j_1, j_2, \ell} |\mathcal{M}^{(\ell)}(j_1, j_2)|$ denotes the maximum sparsity size, $\mathcal{M}^{(\ell)}(j_1, j_2) = \text{supp} \left\{ \beta_0^{(\ell)}(j_1, j_2) \right\}$, and $O(1)$ denotes some positive constant. Similarly, we have $\max_{j, \ell} \|\bar{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 \leq O(1)n^{-1/2}\sqrt{s_0 \log n}$ and $\max_{j, \ell} \|\bar{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_1 \leq O(1)n^{-1/2}s_0\sqrt{\log n}$, with probability approaching one. Therefore, Condition (A5) holds as long as $s_0, s^* = O(n^{\kappa_7})$ for some $\kappa_7 < 1/2$.

Condition (A6) is referred to as the “omega-min” condition in van de Geer and Bühlmann (2013). It essentially guarantees that the true ordering of the mediators can be consistently estimated, which is needed to establish the oracle inequality for the estimator from (11) of Zheng et al. (2018). When the number of mediators d is fixed and the error variances are equal as in (A3), this condition automatically holds.

Finally, we make some remark on why the usual debiasing strategy may not be directly applicable to relax the regularity condition (A4) in Section 5.1 of the paper. Specifically, if (A4) does not hold and the true ordering cannot be recovered, then no matter whether we debias the estimated coefficient matrix or not, the resulting estimator for \mathbf{W}_0 may not be consistent. We illustrate with a simple example. Consider a DAG with two variables, where $X_1 = \varepsilon_1$, $X_2 = aX_1 + \varepsilon_2$ for some $a \neq 0$, and ε_1 and ε_2 are independent mean-zero random errors. Then the corresponding coefficient matrix is

$$\mathbf{W}_0 = \begin{pmatrix} 0 & 0 \\ a & 0 \end{pmatrix}.$$

Meanwhile, we note that the linear structure equation can be rewritten as $X_2 = \varepsilon_2^*$ and $X_1 = a(a^2 + 1)^{-1}X_2 + \varepsilon_1^*$, for some mean-zero random errors ε_1^* and ε_2^* . In addition, with some calculation, we have that

$$\begin{aligned} \varepsilon_2^* &= a\varepsilon_1 + \varepsilon_2, \\ \varepsilon_1^* &= \frac{1}{a^2 + 1}\varepsilon_1 - \frac{a}{a^2 + 1}\varepsilon_2, \end{aligned}$$

and thus ε_1^* and ε_2^* are independent. Then the corresponding coefficient matrix \mathbf{W}_0^* becomes

$$\begin{pmatrix} 0 & a(a^2 + 1)^{-1} \\ 0 & 0 \end{pmatrix}.$$

If the ordering is not correctly specified and (A4) does not hold, then debiasing or not, the second element on the first row will be close to $a(a^2 + 1)^{-1}$, rather than the true value 0. As such, the debiased estimator is not consistent when (A4) is violated. This simple example reflects the challenge of post-selection inference in our setting.

.4 Supporting lemmas

Next, we present two supporting lemmas. Lemma 4 establishes the convergence rate of the variance estimator $\widehat{\sigma}_*^2$, whereas Lemma 5 is needed for the limiting distribution.

Lemma 4. *Suppose (A5) holds, and $\|\mathbf{W}_0\|_2$ is bounded. Then $|\widehat{\sigma}_*^2 - \sigma_*^2| = O(n^{-\kappa_7})$ for some constant κ_7 satisfying $0 < \kappa_7 \leq \min(2\kappa_5, 1/2)$, with probability tending to 1.*

Proof: It suffices to show that, for $\ell = 1, 2$ and any κ_7 satisfying that $\kappa_7 \leq 2\kappa_5$, $\kappa_7 < 1/2$,

$$\frac{1}{(d+2)|\mathcal{I}_\ell^c|} \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell^c} |X_{i,j} - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}_i|_2^2 - \sigma_*^2 = O(n^{-\kappa_7}), \quad (20)$$

with probability tending to 1. In turn, it suffices to show that,

$$\Pr \left(\left| \frac{1}{(d+2)|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \sum_{j=0}^{d+1} |X_{i,j} - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}_i|_2^2 - \sigma_*^2 \right| > 2\kappa_3^2 n^{-\kappa_7} \mid \overline{\mathbf{W}}^{(\ell)} \right) = o_p(1). \quad (21)$$

This is because, if (21) holds, by bounded convergence theorem, we have

$$\Pr \left(\left| \frac{1}{(d+2)|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \sum_{j=0}^{d+1} |X_{i,j} - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}_i|_2^2 - \sigma_*^2 \right| > 2\kappa_3^2 n^{-\kappa_7} \right) = o(1),$$

which in turn yields (20).

For the conditional mean of the left-hand-side of (20) given $\{\mathbf{X}_i : i \in \mathcal{I}_\ell\}$,

$$\begin{aligned}
& \frac{1}{(d+2)} \mathbb{E} \left(\sum_{j=0}^{d+1} |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|_2^2 \mid \overline{\mathbf{W}}^{(\ell)} \right) - \sigma_*^2 \\
&= \frac{1}{(d+2)} \sum_{j=0}^{d+1} \mathbb{E} \|X_j - \mathbf{W}_{0,j}^\top \mathbf{X}\|_2^2 - \sigma_*^2 + \frac{1}{(d+2)} \|\mathbf{W}_0 - \overline{\mathbf{W}}^{(\ell)}\|_2^2 \\
&= \frac{1}{(d+2)} \|\mathbf{W}_0 - \overline{\mathbf{W}}^{(\ell)}\|_2^2 = \frac{1}{(d+2)} \sum_{j=0}^{d+1} \|\mathbf{W}_{0,j} - \overline{\mathbf{W}}_j^{(\ell)}\|_2^2 \leq \kappa_3^2 n^{-2\kappa_5} \leq \kappa_3^2 n^{-\kappa_7},
\end{aligned}$$

where the first equality is due to the fact that $\mathbb{E}\mathbf{X} = 0$ and the second-to-last inequality is due to Condition (A5). The event defined in (21) occurs only when

$$\left| \frac{1}{(d+2)|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \sum_{j=0}^{d+1} |X_{i,j} - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}_i|_2^2 - \frac{1}{(d+2)} \mathbb{E} \left(\sum_{j=0}^{d+1} |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|_2^2 \mid \overline{\mathbf{W}}^{(\ell)} \right) \right| > \kappa_3^2 n^{-\kappa_7}.$$

Thus, to prove (21), it suffices to show that

$$\begin{aligned}
& \Pr \left(\left| \frac{1}{(d+2)|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \sum_{j=0}^{d+1} |X_{i,j} - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}_i|_2^2 - \frac{1}{(d+2)} \mathbb{E} \left(\sum_{j=0}^{d+1} |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|_2^2 \mid \overline{\mathbf{W}}^{(\ell)} \right) \right| \right. \\
& \quad \left. > \kappa_3^2 n^{-\kappa_7} \mid \overline{\mathbf{W}}^{(\ell)} \right) = o_p(1).
\end{aligned}$$

By Chebyshev's inequality, this probability is bounded from above by

$$\frac{n^{2\kappa_7-1}}{2\kappa_3^4(d+2)^2} \text{Var} \left(\sum_{j=0}^{d+1} |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|_2^2 \mid \overline{\mathbf{W}}^{(\ell)} \right) \leq \frac{n^{2\kappa_7-1}}{\kappa_3^4(d+2)^2} \mathbb{E} \left(\left\{ \sum_{j=0}^{d+1} \|X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}\|_2^2 \right\}^2 \mid \overline{\mathbf{W}}^{(\ell)} \right).$$

Since $\kappa_7 < 1/2$, it suffices to show that, with probability approaching one,

$$\frac{1}{(d+2)^2} \mathbb{E} \left(\left\{ \sum_{j=0}^{d+1} |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|_2^2 \right\}^2 \mid \overline{\mathbf{W}}^{(\ell)} \right) = O(1), \tag{22}$$

By Cauchy-Schwarz inequality, the left-hand-side of (22) is bounded from above by

$$\begin{aligned} \frac{1}{d+2} \sum_{j=0}^{d+1} \mathbb{E} \left\{ |X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|^4 \mid \overline{\mathbf{W}}^{(\ell)} \right\} &\leq \\ \frac{16}{d+2} \sum_{j=0}^{d+1} \left(\mathbb{E} \left\{ |X_j|^4 \mid \overline{\mathbf{W}}^{(\ell)} \right\} + \mathbb{E} \left\{ |\overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|^4 \mid \overline{\mathbf{W}}^{(\ell)} \right\} \right). \end{aligned} \quad (23)$$

For any random variable Z and any constant $\kappa > 0$, by Taylor's theorem, we have

$$\mathbb{E} \left\{ \exp \left(\frac{Z^2}{\kappa^2} \right) \right\} = \sum_{k=0}^{+\infty} \frac{\mathbb{E}(Z^{2k})}{\kappa^{2k} k!}.$$

It follows that $\mathbb{E}\{Z^4/(2\kappa^4)\} \leq \mathbb{E}\{\exp(Z^2/\kappa^2)\} - 1$. Let $\|Z\|_{\psi_p}$ denote its Orlicz norm, i.e.,

$$\|Z\|_{\psi_p} = \inf \left[C > 0 : \mathbb{E} \left\{ \exp \left(\frac{|Z|^p}{|C|^p} \right) \leq 2 \right\} \right].$$

By definition, we have $\mathbb{E}\{Z^4/(2\|Z\|_{\psi_2}^4)\} \leq 1$, and henceforth, $\mathbb{E}(Z^4) \leq 2\|Z\|_{\psi_2}^4$.

Under our model assumptions, the covariance matrix $\Sigma_0 = \text{cov}(\mathbf{X})$ is given by $\sigma_*^2(\mathbf{I}_{d+2} - \mathbf{W}_0)^{-1}\{(\mathbf{I}_{d+2} - \mathbf{W}_0)^{-1}\}^\top$, where \mathbf{I}_{d+2} is a $(d+2) \times (d+2)$ identity matrix. For any $\mathbf{a} \in \mathbb{R}^{d+2}$, by Cauchy-Schwarz inequality,

$$\mathbf{a}^\top \Sigma_0^{-1} \mathbf{a} \leq \kappa_* \|\mathbf{a}^\top (\mathbf{I}_{d+2} - \mathbf{W}_0)\|_2^2 \leq \kappa_* \|\mathbf{a}\|_2^2 \|\mathbf{I}_{d+2} - \mathbf{W}_0\|_2^2 \leq 2\kappa_* \|\mathbf{a}\|_2^2 (\|\mathbf{I}_{d+2}\|_2^2 + \|\mathbf{W}_0\|_2^2),$$

for some constant $\kappa_* > 0$. Since $\|\mathbf{W}_0\|_2$ is bounded, it implies that the maximum eigenvalues of Σ_0^{-1} is bounded. Thus, the minimum eigenvalue of Σ_0 is bounded away from zero. Also, note that $(\mathbf{I}_{d+2} - \mathbf{W}_0)^{-1} = \mathbf{I}_{d+2} + \mathbf{W}_0$. Following similar arguments, the maximum eigenvalue of Σ_0 is bounded as well. Since \mathbf{X} is jointly normal with bounded $\lambda_{\max}(\Sigma_0)$, we have, for some constant $\kappa > 0$ and any vector $\mathbf{a} \in \mathbb{R}^{d+2}$,

$$\|\mathbf{a}^\top \mathbf{X}\|_{\psi_2} \leq \kappa \|\mathbf{a}\|_2. \quad (24)$$

It then follows from (24) that conditional on $\overline{\mathbf{W}}^{(\ell)}$,

$$\mathbb{E}(|X_j|^4) \leq 2\|X_j\|_{\psi_2}^4 \leq 2\kappa^4, \quad \text{and} \quad \mathbb{E}(|\overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|^4) \leq 2\|\overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}\|_{\psi_2}^4 \leq 2\kappa^4 \|\overline{\mathbf{W}}_j^{(\ell)}\|_2^4,$$

for any $j = 1, \dots, d$. It follows from (23) that conditional on $\overline{\mathbf{W}}^{(\ell)}$,

$$\frac{1}{(d+2)} \sum_{j=1}^d \mathbb{E}(|X_j - \overline{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|^4) \leq \frac{32\kappa^4}{d+2} \sum_{j=0}^{d+1} \left(1 + \|\overline{\mathbf{W}}_j^{(\ell)}\|_2^4\right) \leq 32\kappa^4 \left(1 + \max_j \|\overline{\mathbf{W}}_j^{(\ell)}\|_2^4\right).$$

By (A5), we have $\max_j \|\bar{\mathbf{W}}_j^{(\ell)}\|_2 \leq \max_j \|\mathbf{W}_{0,j}\|_2 + \max_j \|\mathbf{W}_{0,j} - \bar{\mathbf{W}}_j^{(\ell)}\|_2 \leq \max_j \|\mathbf{W}_{0,j}\|_2 + \kappa_3 n^{-\kappa_5}$. Since $\|\mathbf{W}_0\|_2$ is bounded, we have $\|\mathbf{e}_j^\top \mathbf{W}_0\|_2 = O(1)$, and hence $\|\mathbf{W}_{0,j}\|_2 = O(1)$, where \mathbf{e}_j denotes a $(d+2)$ -dimensional vector with the j th element equal to one and the rest equal to zero. It follows that $\sum_{j=0}^{d+1} \mathbb{E} \left\{ |X_j - \bar{\mathbf{W}}_j^{(\ell)\top} \mathbf{X}|^4 / (d+2) \mid \bar{\mathbf{W}}^{(\ell)} \right\} = O(1)$, with probability approaching one. Then (22) is proven. This completes the proof. \square

Lemma 5. *Suppose (A1), (A4) hold. Then $\text{ACT}(j, \widetilde{\mathbf{W}}^{(\ell)})$ contains no descendants of j .*

Proof: Suppose there exists some $j' \in \text{ACT}(j, \widetilde{\mathbf{W}}^{(\ell)})$, such that j' is a descendant of j . By definition, there exists a directed path from X_j to $X_{j'}$: $X_j \rightarrow X_{i_1} \rightarrow \dots \rightarrow X_{i_K} \rightarrow X_{j'}$. By Condition (A4), we have $j \in \text{ACT}(i_1, \widetilde{\mathbf{W}}^{(\ell)})$, $i_k \in \text{ACT}(i_{k+1}, \widetilde{\mathbf{W}}^{(\ell)})$, for $k = 1, \dots, K-1$ and $i_K \in \text{ACT}(j', \widetilde{\mathbf{W}}^{(\ell)})$. This, together with $j' \in \text{ACT}(j, \widetilde{\mathbf{W}}^{(\ell)})$, implies that there exists a directed path from X_j to X_j on the DAG generated by $\widetilde{\mathbf{W}}^{(\ell)}$. Then the acyclic constraint of $\widetilde{\mathbf{W}}^{(\ell)}$ is violated. This completes the proof. \square

.5 Proof of Lemma 2

To prove Lemma 2, it suffices to show that

$$(|\mathbf{W}_0|^{(k)})_{q_2, q_1} = \max_{0 \leq j_1, \dots, j_{k-1} \leq d+1} \min \left(|W_{0,j_1, q_1}|, \min_{l \in \{1, \dots, k-2\}} |W_{0,j_{l+1}, j_l}|, |W_{0,q_2, j_{k-1}}| \right). \quad (25)$$

Lemma 2 can then be similarly proven as Lemma 1. We use induction to prove (25) for any $q_1 = 0, \dots, d$, and $q_2 = 1, \dots, d+1$. When $k = 2$, by the definition of \otimes , we have $(|\mathbf{W}_0|^{(2)})_{q_2, q_1} = \max_{0 \leq j \leq d+1} \min(|W_{0,j, q_1}|, |W_{0,q_2, j}|)$. Thus, (25) holds with $k = 2$.

Suppose (25) holds with $k = t$ for some $t \geq 2$, i.e.,

$$(|\mathbf{W}_0|^{(t)})_{q_2, q_1} = \max_{0 \leq j_1, \dots, j_{t-1} \leq d+1} \min \left(|W_{0,j_1, q_1}|, \min_{l \in \{1, \dots, t-2\}} |W_{0,j_{l+1}, j_l}|, |W_{0,q_2, j_{t-1}}| \right). \quad (26)$$

Therefore,

$$\begin{aligned} (|\mathbf{W}_0|^{(t+1)})_{q_2, q_1} &= (|\mathbf{W}_0|^{(t)} \circ |\mathbf{W}_0|)_{q_2, q_1} = \max_{j \in \{0, \dots, d+1\}} \min\{(|\mathbf{W}_0|^{(t)})_{j, q_1}, |W_{0,q_2, j}|\} \\ &= \max_{j \in \{0, \dots, d+1\}} \min \left\{ \max_{0 \leq j_1, \dots, j_{t-1} \leq d+1} \min \left(|W_{0,j_1, q_1}|, \min_{l \in \{1, \dots, t-2\}} |W_{0,j_{l+1}, j_l}|, |W_{0,q_2, j_{t-1}}| \right), |W_{0,q_2, j}| \right\} \\ &= \max_{j \in \{0, \dots, d+1\}} \max_{0 \leq j_1, \dots, j_{t-1} \leq d+1} \min \left(|W_{0,j_1, q_1}|, \min_{l \in \{1, \dots, t-2\}} |W_{0,j_{l+1}, j_l}|, |W_{0,q_2, j_{t-1}}|, |W_{0,q_2, j}| \right) \\ &= \max_{0 \leq j_1, \dots, j_{t-1}, j_t \leq d+1} \min \left(|W_{0,j_1, q_1}|, \min_{l \in \{1, \dots, t-1\}} |W_{0,j_{l+1}, j_l}|, |W_{0,q_2, j_t}| \right). \end{aligned}$$

Thus, (26) holds with $k = t + 1$. This completes the proof of Lemma 2. \square

.6 Proof of Theorem 1

Outline of the proof: Our goal is to prove

$$\Pr \left\{ \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{i,j}^{(\ell)} - W_{0,i,j}| > \widehat{c}^{(\ell)}(q_1, q_2) \right\} = \frac{\alpha}{2} + o(1), \quad (27)$$

for any $q_1 = 0, \dots, d, q_2 = 1, \dots, d + 1$ and $\ell = 1, 2$. Then the validity of our test follows by the union-intersection principle. We begin with an outline of our proof, which relies on the high-dimensional central limit theorem that was recently developed by Chernozhukov et al. (2013b) and Chernozhukov et al. (2014). We divide the proof into three steps.

In Step 1, we show that,

$$\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| \sqrt{|\mathcal{I}_\ell^c|} \left(\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2} \right) - \eta_{0, j_1, j_2}^{(\ell)} \right| = o_p(\log^{-1/2} n), \quad (28)$$

where

$$\eta_{0, j_1, j_2}^{(\ell)} = \frac{|\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \left\{ X_{i, j_2} - \beta_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} \varepsilon_{i, j_1}}{\mathbb{E} \left[X_{j_1} \left\{ X_{j_2} - \beta_0^{(\ell)\top}(j_1, j_2) \mathbf{X} \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right]}. \quad (29)$$

This further implies that

$$\underbrace{\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \sqrt{|\mathcal{I}_\ell^c|} |\widehat{W}_{j_1, j_2}^{(\ell)} - W_{0, j_1, j_2}|}_{\widehat{S}^{(\ell)}} = \underbrace{\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{0, j_1, j_2}^{(\ell)}|}_{S_0^{(\ell)}} + o_p(\log^{-1/2} n). \quad (30)$$

Let $\mathcal{M}^{(\ell)}(j_1, j_2)$ denote the support of $\beta_0^{(\ell)}(j_1, j_2)$. By definition, for any $j \in \mathcal{M}^{(\ell)}(j_1, j_2)$, we have $j, j_2 \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})$. Since $\widetilde{\mathbf{W}}^{(1)}$ and $\widetilde{\mathbf{W}}^{(2)}$ satisfy the acyclic constraint in (A1), by Lemma 5, neither j nor j_2 is a descendant of j_1 . As a result, X_{j_1} is conditionally independent of X_{j_2} and X_j given its parents. As such, the numerator of $\eta_{0, j_1, j_2}^{(\ell)}$ forms a sum of i.i.d. mean zero random variables.

In Step 2, we show that,

$$\sup_{z \in \mathbb{R}} \left| \Pr \left(S_0^{(\ell)} \leq z \mid \widetilde{\mathbf{W}}^{(\ell)} \right) - \Pr \left(\|N(0, V_0)\|_\infty \leq z \mid \widetilde{\mathbf{W}}^{(\ell)} \right) \right| = o(1), \quad (31)$$

where V_0 is a matrix involving the covariance of $\eta_{0, j_1, j_2}^{(\ell)}$ in (29) and is defined in Step 2.

In Step 3, we show that, for some constant $\kappa^* > 0$,

$$\|V_0 - \widehat{V}\|_{\infty, \infty} = O_p(n^{-\kappa^*}), \quad (32)$$

where $\|\cdot\|_{\infty, \infty}$ denotes the elementwise max-norm, and \widehat{V} is a matrix involving the covariance of $\eta_{j_1, j_2}^{*(\ell)}$ in (16) of the paper, and is defined later in Step 3.

Following (31) and bounded convergence theorem, we have that,

$$\sup_{z \in \mathbb{R}} \left| \Pr \left(S_0^{(\ell)} \leq z \right) - \Pr \left(\|N(0, V_0)\|_{\infty} \leq z \right) \right| = o(1).$$

This, together with (30), yields that

$$\begin{aligned} \Pr \left(\widehat{S}^{(\ell)} \leq z \right) &\geq \Pr \left(\|N(0, V_0)\|_{\infty} \leq z - \varepsilon \log^{-1/2} n \right) - o(1), \\ \Pr \left(\widehat{S}^{(\ell)} \leq z \right) &\leq \Pr \left(\|N(0, V_0)\|_{\infty} \leq z + \varepsilon \log^{-1/2} n \right) + o(1), \end{aligned} \quad (33)$$

for any sufficiently small $\varepsilon > 0$, where the little-o term is uniform in z . Using similar arguments for (33) and also Lemma 3.1 of Chernozhukov et al. (2015), we have by (32)

$$\begin{aligned} \Pr \left(\widehat{S}^{(\ell)} \leq z \right) &\geq \Pr \left(\|N(0, \widehat{V})\|_{\infty} \leq z - 2\varepsilon \log^{-1/2} n | \widehat{V} \right) - o(1), \\ \Pr \left(\widehat{S}^{(\ell)} \leq z \right) &\leq \Pr \left(\|N(0, \widehat{V})\|_{\infty} \leq z + 2\varepsilon \log^{-1/2} n | \widehat{V} \right) + o(1), \end{aligned}$$

for any sufficiently small $\varepsilon > 0$. Set $z = \widehat{c}^{(\ell)}(q_1, q_2)$. Since the little-o term is uniform in $z \in \mathbb{R}$, we have

$$\begin{aligned} \Pr \left\{ \widehat{S}^{(\ell)} \leq \widehat{c}^{(\ell)}(q_1, q_2) \right\} &\geq \Pr \left\{ \|N(0, \widehat{V})\|_{\infty} \leq \widehat{c}^{(\ell)}(q_1, q_2) - 2\varepsilon \log^{-1/2} n | \widehat{V} \right\} - o(1), \\ \Pr \left\{ \widehat{S}^{(\ell)} \leq \widehat{c}^{(\ell)}(q_1, q_2) \right\} &\leq \Pr \left\{ \|N(0, \widehat{V})\|_{\infty} \leq \widehat{c}^{(\ell)}(q_1, q_2) + 2\varepsilon \log^{-1/2} n | \widehat{V} \right\} + o(1). \end{aligned} \quad (34)$$

We show in Step 2 that all diagonal elements in V_0 are well bounded away from zero. Henceforth, with probability approaching 1, all diagonal elements in \widehat{V} are well bounded away from zero as well. It follows from Theorem 1 of Chernozhukov et al. (2017) that

$$\begin{aligned} \Pr \left\{ \|N(0, \widehat{V})\|_{\infty} \leq \widehat{c}^{(\ell)}(q_1, q_2) + 2\varepsilon \log^{-1/2} n | \widehat{V} \right\} &- \Pr \left\{ \|N(0, \widehat{V})\|_{\infty} \leq \widehat{c}^{(\ell)}(q_1, q_2) \right. \\ &\left. - 2\varepsilon \log^{-1/2} n | \widehat{V} \right\} \leq O(1)\varepsilon \log^{1/2} d \log^{-1/2} n, \end{aligned}$$

where $O(1)$ denotes some positive constant. Under the given conditions, we have $\log d = O(\log n)$. The right-hand-side is bounded by $\kappa\varepsilon$ for some constant $\kappa > 0$. This, together

with (34), yields that

$$\left| \Pr \left\{ \widehat{S}^{(\ell)} \leq \widehat{c}^{(\ell)}(q_1, q_2) \right\} - \Pr \left\{ \|N(0, \widehat{V})\|_\infty \leq \widehat{c}^{(\ell)}(q_1, q_2) | \widehat{V} \right\} \right| \leq \kappa \varepsilon + o(1).$$

Since ε can be made arbitrarily small, (27) follows, which completes the proof of Theorem 1. Next, we give detailed proofs for each step.

Step 1: The proof of this step relies on the arguments developed to establish the limiting distribution of the debiased LASSO (van de Geer et al., 2014) and the decorrelated score statistic (Ning and Liu, 2017). We aim to establish the upper bound for $\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\sqrt{|\mathcal{I}_\ell^c|} (\widehat{W}_{j_1, j_2} - W_{0, j_1, j_2}) - \eta_{j_1, j_2}^{(\ell)}|$, and for $\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell)} - \eta_{0, j_1, j_2}^{(\ell)}|$. Together, these two upper bounds would lead to (28).

To obtain the first upper bound, we have that,

$$\begin{aligned} & \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| \sqrt{|\mathcal{I}_\ell^c|} \left(\widehat{W}_{j_1, j_2} - W_{0, j_1, j_2} \right) - \eta_{j_1, j_2}^{(\ell)} \right| \\ & \leq \left| \frac{|\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \left\{ X_{i, j_2} - \widehat{\beta}^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} \left\{ \sum_{j \neq j_2} X_{i, j} (\overline{W}_{j_1, j}^{(\ell)} - W_{0, j_1, j}) \right\}}{|\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} X_{i, j_2} \left\{ X_{i, j_2} - \widehat{\beta}^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\}} \right| = \frac{I_1(j_1, j_2, \ell)}{I_2(j_1, j_2, \ell)}. \end{aligned}$$

We next bound $I_1(j_1, j_2, \ell)$ and $I_2(j_1, j_2, \ell)$, respectively.

For $I_1(j_1, j_2, \ell)$, we have that, $I_1(j_1, j_2, \ell) \leq I_{1,1}(j_1, j_2, \ell) + I_{1,2}(j_1, j_2, \ell)$, where

$$\begin{aligned} I_{1,1}(j_1, j_2, \ell) &= \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \beta_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i - \widehat{\beta}^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} \left\{ \sum_{j \neq j_2} X_{i, j} (\overline{W}_{j_1, j}^{(\ell)} - W_{0, j_1, j}) \right\} \right|, \\ I_{1,2}(j_1, j_2, \ell) &= \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \left\{ X_{i, j_1} - \beta_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} \left\{ \sum_{j \neq j_2} X_{i, j} (\overline{W}_{j_1, j}^{(\ell)} - W_{0, j_1, j}) \right\} \right|. \end{aligned}$$

For $I_{1,1}(j_1, j_2)$, by Cauchy-Schwarz inequality, it can be upper bounded by

$$\begin{aligned} & |\mathcal{I}_\ell^c|^{1/2} \|\widehat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^\ell(j_1, j_2)\|_2 \|\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0, j_1}\|_2 \|\Sigma_0\|_2 \\ & + |\mathcal{I}_\ell^c|^{1/2} \|\widehat{\beta}^{(\ell)}(j_1, j_2) - \beta_0^\ell(j_1, j_2)\|_1 \|\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0, j_1}\|_1 \left\| |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \mathbf{X}_i \mathbf{X}_i^\top - \Sigma_0 \right\|_{\infty, \infty}, \end{aligned}$$

where $\|\cdot\|_{\infty, \infty}$ denotes the elementwise maximum norm in absolute value. In the proof of Lemma 4, we have shown that $\|\Sigma_0\|_2$ is bounded. By Condition (A5), the first line is upper

bounded by $O(1)\sqrt{nn}^{-(\kappa_4+\kappa_5)}$ where $O(1)$ denotes some positive constant.

In addition, using similar arguments as in Equation (A.70) of Shi et al. (2020), we have

$$\left\| |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \mathbf{X}_i \mathbf{X}_i^\top - \boldsymbol{\Sigma}_0 \right\|_{\infty, \infty} = O(n^{-1/2} \sqrt{\log d}),$$

with probability approaching 1. This together with the condition on the number of mediators d and Condition (A5) implies that the second line is $O(n^{-\kappa_6} \sqrt{\log n})$, with probability approaching 1. Consequently,

$$\max_{j_1, j_2, \ell} I_{1,1}(j_1, j_2, \ell) \leq O(1) n^{\max(1/2 - (\kappa_4 + \kappa_5), -\kappa_6)} \sqrt{\log n}, \quad (35)$$

where $O(1)$ denotes some positive constant.

For $I_{1,2}(j_1, j_2)$, note that $W_{0,j_1,j} \neq 0$ only when j is a parent of j_1 . By (A4), we have $W_{0,j_1,j} \neq 0$ only when $j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})$. It follows that,

$$\sum_{j \neq j_2} X_{i,j} \left(\overline{W}_{j_1,j}^{(\ell)} - W_{0,j_1,j} \right) = \sum_{j \neq j_2, j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})} X_{i,j} \left(\overline{W}_{j_1,j}^{(\ell)} - W_{0,j_1,j} \right).$$

By the definition of $\beta_0^{(\ell)}(j_1, j_2)$, we have $\mathbb{E} \left[\left\{ X_{j_2} - \mathbf{X}^\top \beta_0^{(\ell)}(j_1, j_2) \right\} X_j | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)} \right] = 0$ for any $j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)}) - \{j_2\}$. Note that $\widetilde{\mathbf{W}}^{(\ell)}$ and $\overline{\mathbf{W}}^{(\ell)}$ are constructed by samples in $\{X_{i,j}\}_{i \in \mathcal{I}_\ell}$. It follows that $\mathbb{E} \left\{ \varphi_i^{(\ell)}(j_1, j_2) | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)} \right\} = 0$ for any $i \in \mathcal{I}_\ell$, where

$$\varphi_i^{(\ell)}(j_1, j_2) = \left\{ X_{i,j_2} - \mathbf{X}_i^\top \beta_0^{(\ell)}(j_1, j_2) \right\} \left\{ \sum_{j \neq j_2, j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})} X_{i,j} (\overline{W}_{j_1,j}^{(\ell)} - W_{0,j_1,j}) \right\}.$$

By the definition of the Orlicz norm and Cauchy-Schwarz inequality, we have

$$\begin{aligned} \|\varphi_i^{(\ell)}(j_1, j_2)\|_{\psi_1 | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}} &= \|\varphi_i^{(\ell)}(j_1, j_2)\|_{\psi_2 | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}}^2 \\ &\leq \frac{1}{2\tau} \|X_{i,j_2} - \mathbf{X}_i^\top \beta_0^{(\ell)}(j_1, j_2)\|_{\psi_2 | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}}^2 + \frac{\tau}{2} \|\mathbf{X}_i^\top (\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0,j_1})\|_{\psi_2 | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}}^2, \end{aligned} \quad (36)$$

where $\|\cdot\|_{\psi_p | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}}$ denotes the Orlicz norm conditional on $\widetilde{\mathbf{W}}^{(\ell)}$ and $\overline{\mathbf{W}}^{(\ell)}$. Since j_2 does not belong to the support of $\beta_0^{(\ell)}(j_1, j_2)$, we have by (24) that

$$\|X_{i,j_2} - \mathbf{X}_i^\top \beta_0^{(\ell)}(j_1, j_2)\|_{\psi_2 | \widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}} \leq c_0 \left\{ 1 + \|\beta_0^{(\ell)}(j_1, j_2)\|_2 \right\}.$$

Using similar arguments in Lemma A.1 of Shi et al. (2020), we have $\max_{j_1, j_2, \ell} \|\beta_0^{(\ell)}(j_1, j_2)\|_2 =$

$O(1)$. Henceforth, $\max_{i,j_1,j_2,\ell} \|X_{i,j_2} - \mathbf{X}_i^\top \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2)\|_{\psi_2|\widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}} = O(1)$. Similarly, we have,

$$\max_{i,j_1,j_2,\ell} \|\mathbf{X}_i^\top (\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0,j_1})\|_{\psi_2|\widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}} \leq c_0 \max_{j_1} \|\overline{\mathbf{W}}_{j_1}^{(\ell)} - \mathbf{W}_{0,j_1}\|_2 \leq c_0 \kappa_3 n^{-\kappa_5},$$

by Condition (A5). Setting $\tau = n^{-\kappa_5}$, it follows from (36) that

$$\max_{i,\ell} \|\varphi_i^{(\ell)}(j_1, j_2)\|_{\psi_1|\widetilde{\mathbf{W}}^{(\ell)}, \overline{\mathbf{W}}^{(\ell)}} = O(n^{-\kappa_5}), \quad (37)$$

by (A4) and (A5). It follows from Lemma G.3 of Shi et al. (2018) that

$$\Pr \left(\left| \sum_{i \in \mathcal{I}_\ell^c} \varphi_i^{(\ell)}(j_1, j_2) \right| > t \mid \max_j \|\overline{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 \leq \kappa_3 n^{-\kappa_5}, \widetilde{\mathbf{W}}^{(\ell)} \right) \leq 2 \exp \left\{ -\kappa \min \left(\frac{t^2}{n^{1-2\kappa_5}}, \frac{t}{n^{-\kappa_5}} \right) \right\},$$

for some constant $\kappa > 0$. By Bonferroni's inequality, we have that,

$$\begin{aligned} & \Pr \left(\max_{j_1, j_2, \ell} \left| \sum_{i \in \mathcal{I}_\ell^c} \varphi_i^{(\ell)}(j_1, j_2) \right| > t \mid \max_j \|\overline{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 \leq \kappa_3 n^{-\kappa_5}, \widetilde{\mathbf{W}}^{(\ell)} \right) \\ & \leq 4(d+2)^2 \exp \left\{ -c \min \left(\frac{t^2}{n^{1-2\kappa_5}}, \frac{t}{n^{-\kappa_5}} \right) \right\} \leq 4 \exp \left\{ -c \min \left(\frac{t^2}{n^{1-2\kappa_5}}, \frac{t}{n^{-\kappa_5}} \right) + 2 \log(d+2) \right\}. \end{aligned}$$

Setting $t = 3\kappa_1 c^{-1} n^{1/2-\kappa_5} \sqrt{\log n}$, for a sufficiently large n , we have that,

$$\begin{aligned} 4 \exp \left\{ -c \min \left(\frac{t^2}{n^{1-2\kappa_5}}, \frac{t}{n^{-\kappa_5}} \right) + 2 \log(d+2) \right\} &= 4 \exp \{ -3\kappa_1 \log n + 2 \log(d+2) \} \\ &= \frac{4(d+2)^2}{n^{3\kappa_1}} = O(n^{-\kappa_1}) = o(1). \end{aligned}$$

Therefore, conditional on the events in (A4) and (A5), we have that, with probability approaching one, $\max_{j_1, j_2, \ell} |\sum_{i \in \mathcal{I}_\ell^c} \varphi_i^{(\ell)}(j_1, j_2)| = O(n^{1/2-\kappa_5} \sqrt{\log n})$, or equivalently,

$$\max_{j_1, j_2, \ell} I_{1,2}(j_1, j_2, \ell) = O(n^{-\kappa_5} \sqrt{\log n}). \quad (38)$$

Combining (35) and (38) together yields, for some constant $\kappa > 0$,

$$\max_{j_1, j_2, \ell} I_1(j_1, j_2, \ell) = O(n^{-\kappa}), \quad (39)$$

For $I_2(j_1, j_2, \ell)$, we first define

$$I_2^*(j_1, j_2, \ell) = |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \left\{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} X_{i,j_2},$$

$$I_2^{**}(j_1, j_2, \ell) = \mathbb{E} \left\{ I_2^*(j_1, j_2, \ell) \mid \widetilde{\mathbf{W}}^{(\ell)} \right\} = \mathbb{E} \left[\left\{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right].$$

Following the proof of Corollaries 4.1 and 4.2 of Ning and Liu (2017), we have that,

$$\begin{aligned}
|I_2(j_1, j_2, \ell) - I_2^*(j_1, j_2, \ell)| &\leq |\mathcal{I}_\ell^c|^{-1} \left| \sum_{i \in \mathcal{I}_\ell^c} X_{i,j_2} \{ \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2) - \widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) \}^\top \mathbf{X}_i \right| \\
&\leq \| \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2) - \widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) \|_1 \left(\frac{2}{n} \left\| \sum_{i \in \mathcal{I}_\ell^c} X_{i,j_2} \mathbf{X}_i - |\mathcal{I}_\ell^c| \mathbb{E} X_{j_2} \mathbf{X} \right\|_\infty \right) \\
&+ \| \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2) - \widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) \|_2 \lambda_{\max}(\boldsymbol{\Sigma}_0) = O(n^{-\kappa_4}) + O(n^{-\kappa_6-1/2} \sqrt{\log n}), \tag{40}
\end{aligned}$$

where the big-O term is uniform in (j_1, j_2, ℓ) . In addition, note that

$$\begin{aligned}
\mathbb{E} |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \left[X_{i,j_2} \left\{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} | \widetilde{\mathbf{W}}^{(\ell)} \right] &= \mathbb{E} \left[X_{j_2} \left\{ X_{j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X} \right\} | \widetilde{\mathbf{W}}^{(\ell)} \right] \\
&= \mathbb{E} \left[\left\{ X_{j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X} \right\}^2 | \widetilde{\mathbf{W}}^{(\ell)} \right] \geq \| \{1, \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)\} \|_2^2 \lambda_{\min}(\boldsymbol{\Sigma}_0) \geq \lambda_{\min}(\boldsymbol{\Sigma}_0).
\end{aligned}$$

Since the minimum eigenvalue of $\boldsymbol{\Sigma}_0$ is bounded away from zero, we have

$$\min_{j_1, j_2, \ell} I_2^{**}(j_1, j_2, \ell) = \min_{j_1, j_2, \ell} \mathbb{E} \left\{ I_2^*(j_1, j_2, \ell) | \widetilde{\mathbf{W}}^{(\ell)} \right\} \geq 2\varepsilon, \tag{41}$$

for some $\varepsilon > 0$. Similar to (38), we have,

$$\max_{j_1, j_2, \ell} \left\| \sum_{i \in \mathcal{I}_\ell^c} \{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \} X_{i,j_2} - |\mathcal{I}_\ell^c| I_2^{**}(j_1, j_2, \ell) \right\|_\infty = O(\sqrt{n \log n}). \tag{42}$$

This together with (41) yields,

$$\min_{j_1, j_2, \ell} I_2^*(j_1, j_2, \ell) = \min_{j_1, j_2, \ell} |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \} X_{i,j_2} \geq 2\varepsilon.$$

Combining this together with (40), we have,

$$\begin{aligned}
\min_{j_1, j_2, \ell} I_2(j_1, j_2, \ell) &\geq \min_{j_1, j_2, \ell} \left[|\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \left\{ X_{i,j_2} - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} X_{i,j_2} \right. \\
&\quad \left. + |\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} X_{i,j_2} \left\{ \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2) - \widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) \right\}^\top \mathbf{X}_i \right] \geq \varepsilon. \tag{43}
\end{aligned}$$

Combining (39) and (43) together yields,

$$\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| \sqrt{|\mathcal{I}_\ell^c|} (\widehat{W}_{j_1, j_2} - W_{0, j_1, j_2}) - \eta_{j_1, j_2}^{(\ell)} \right| \leq O(1) n^{-\kappa}, \tag{44}$$

for some constant $\kappa > 0$. This gives the first desired upper bound.

To obtain the second upper bound, we have that,

$$\begin{aligned} \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell)} - \eta_{0, j_1, j_2}^{(\ell)}| &\leq \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \frac{|\mathcal{I}_\ell^c|^{-1/2} |\sum_{i \in \mathcal{I}_\ell^c} \{\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2)\}^\top \mathbf{X}_{i\varepsilon_{i, j_1}}|}{I_2(j_1, j_2, \ell)} \\ &+ \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \frac{|I_2(j_1, j_2, \ell) - I_2^{**}(j_1, j_2, \ell)| |\mathcal{I}_\ell^c|^{-1/2} |\sum_{i \in \mathcal{I}_\ell^c} \{\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)}(j_1, j_2)\}^\top \mathbf{X}_{i\varepsilon_{i, j_1}}|}{I_2(j_1, j_2, \ell) I_2^{**}(j_1, j_2, \ell)}. \end{aligned}$$

By (40), (41), (42) and (43), we have,

$$\begin{aligned} &\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell)} - \eta_{0, j_1, j_2}^{(\ell)}| \\ &\leq \frac{2}{\varepsilon} \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \{\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)\}^\top \mathbf{X}_{i\varepsilon_{i, j_1}} \right|. \end{aligned} \quad (45)$$

Note that the supports of $\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2)$ and $\boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)$ belong to $\text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})$. Since ε_{i, j_1} is independent of $\{X_{i, j} : j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})\}$ conditional on $\widetilde{\mathbf{W}}^{(\ell)}$, we have,

$$\begin{aligned} &\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \{\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)\}^\top \mathbf{X}_{i\varepsilon_{i, j_1}} \right| \\ &\leq \max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \|\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)\|_1 \max_{j \in \text{ACT}(j_1, \widehat{\mathbf{B}}^{(\ell)})} \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} X_{i, j} \varepsilon_{i, j_1} \right|. \end{aligned}$$

By (A4), ε_{i, j_1} is conditionally independent of X_j for any $j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})$. Similar to (38), by Bernstein's inequality, we have that,

$$\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \max_{j \in \text{ACT}(j_1, \widetilde{\mathbf{W}}^{(\ell)})} \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} X_{i, j} \varepsilon_{i, j_1} \right| = O(\sqrt{\log n}).$$

It follows from Condition (A5) that,

$$\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} \left| |\mathcal{I}_\ell^c|^{-1/2} \sum_{i \in \mathcal{I}_\ell^c} \{\widehat{\boldsymbol{\beta}}^{(\ell)}(j_1, j_2) - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, j_2)\}^\top \mathbf{X}_{i\varepsilon_{i, j_1}} \right| = O(n^{-\kappa}),$$

for some constant $\kappa > 0$. In view of (45), we obtain the second desired upper bound,

$$\max_{(j_1, j_2) \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})} |\eta_{j_1, j_2}^{(\ell)} - \eta_{0, j_1, j_2}^{(\ell)}| = O(n^{-\kappa}), \quad (46)$$

for some constant $\kappa > 0$.

Finally, combining (44) and (46) together yields (28). This completes Step 1.

Step 2: Recall that, each ε_{i,j_1} is uncorrelated with $X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i$. Since both ε_{i,j_1} and $X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i$ are normally distributed given $\widetilde{\mathbf{W}}^{(\ell)}$, they are independent as well. As a result, we have,

$$\text{Var} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\} \varepsilon_{i,j_1} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] = \sigma_*^2 \text{E} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right],$$

for $1 \leq j_1 \leq d+1$. Recall that $I_2^{**}(j_1, j_2, \ell) = \text{E} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right]$. Then, we have for $1 \leq j_1 \leq d+1$ that

$$\begin{aligned} \text{Var} \left(\eta_{0,j_1,j_2}^{(\ell)} \mid \widetilde{\mathbf{W}}^{(\ell)} \right) &= \frac{1}{\{I_2^{**}(j_1, j_2, \ell)\}^2} \text{Var} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\} \varepsilon_{i,j_1} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \\ &= \frac{1}{\sigma_*^2 \text{E} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right]}. \end{aligned} \quad (47)$$

Since $j_2 \notin \mathcal{M}^{(\ell)}(j_1, j_2)$, we have

$$\begin{aligned} \lambda_{\min}(\Sigma_0) \left\| \{1, \beta_0^{(\ell)\top}(j_1, j_2)\}^\top \right\|_2^2 &\leq \text{E} \left[\left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \\ &\leq \lambda_{\max}(\Sigma_0) \left\| (1, \beta_0^{(\ell)\top}(j_1, j_2))^\top \right\|_2^2. \end{aligned}$$

Following the proof of Lemma A.1 of Shi et al. (2020), we have $\max_{j_1, j_2, \ell} \|\beta_0^{(\ell)}(j_1, j_2)\|_2 = O(1)$. Since $\lambda_{\min}(\Sigma_0)$ and $\lambda_{\max}(\Sigma_0)$ are uniformly bounded away from zero and infinity, there exists some constant $\kappa \geq 1$ such that, for any j_1, j_2 ,

$$\kappa^{-1} \leq \text{E}[\{X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2)\mathbf{X}_i\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)}] \leq \kappa. \quad (48)$$

It follows from (47) that, for any j_1, j_2 ,

$$\kappa^{-1} \leq \text{Var}(\eta_{0,j_1,j_2}^{(\ell)} \mid \widetilde{\mathbf{W}}^{(\ell)}) \leq \kappa. \quad (49)$$

We index all pairs of indices (j_1, j_2) in $\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})$ by $\{j_1(0), j_2(0)\}, \{j_1(1), j_2(1)\}, \dots, \{j_1(L-1), j_2(L-1)\}$, where $L = |\mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})|$. Next, define a covariance matrix $V_0 \in \mathbb{R}^{L \times L}$, such that its (l_1, l_2) th entry is the covariance of $\eta_{0,j_1(l_1),j_2(l_1)}^{(\ell)}$ and $\eta_{0,j_1(l_2),j_2(l_2)}^{(\ell)}$ conditional on $\widetilde{\mathbf{W}}^{(\ell)}$. By (49), the diagonal elements of V_0 are uniformly bounded away from zero and infinity. We also comment that $\eta_{j_1,j_2}^{(\ell)}$ defined in (15) in Section 4.4 of the paper can be viewed as a more intuitive version of $\eta_{0,j_1,j_2}^{(\ell)}$ defined in (29).

When L is finite, by the classical Lindeberg-Feller central limit theorem and Condition

(A4), we have (31) holds. When L diverges, we have $L \leq d^2$ that grows at an polynomial order of n . By (49), we have $\min_{j_1, j_2, \ell} \text{Var} \left(\eta_{0, j_1, j_2}^{(\ell)} | \widetilde{\mathbf{W}}^{(\ell)} \right)$ is uniformly bounded away from zero. Moreover, following similar arguments in proving (37), we have $\max_{j_1, j_2, \ell} \|\eta_{0, j_1, j_2}^{(\ell)}\|_{\psi_1 | \widetilde{\mathbf{W}}^{(\ell)}} = O(1)$. Then, by Corollary 2.1 of Chernozhukov et al. (2013b), (31) holds as well. This completes Step 2.

Step 3: Define a covariance matrix $\widehat{V} \in \mathbb{R}^{L \times L}$, such that its (l_1, l_2) th entry is the covariance of $\eta_{j_1(l_1), j_2(l_1)}^{*(\ell)}$ and $\eta_{j_1(l_2), j_2(l_2)}^{*(\ell)}$ conditional on the data. To bound $\|\widehat{V} - V_0\|_{\infty, \infty}$ in (32), it suffices to bound

$$\max_{\substack{(j_1, j_2), (j_3, j_4) \\ \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})}} \left| \text{cov} \left[\eta_{j_1, j_2}^{*(\ell)}, \eta_{j_3, j_4}^{*(\ell)} | \{X_{i,j}\}_{1 \leq i \leq n, 0 \leq j \leq d+1} \right] - \text{cov} \left[\eta_{0, j_1, j_2}^{(\ell)}, \eta_{0, j_3, j_4}^{(\ell)} | \{X_{i,j}\}_{i \in \mathcal{I}_\ell^c, 0 \leq j \leq d+1} \right] \right|.$$

Recall in Section 4.4, when $j_1 \neq j_3$, we have $\text{cov}(\eta_{j_1, j_2}^{*(\ell)}, \eta_{j_3, j_4}^{*(\ell)} | \{X_{i,j}\}_{1 \leq i \leq n, 0 \leq j \leq d+1}) = 0$. Similarly, $\text{cov}(\eta_{0, j_1, j_2}^{(\ell)}, \eta_{0, j_3, j_4}^{(\ell)} | \{X_{i,j}\}_{i \in \mathcal{I}_\ell^c, 0 \leq j \leq d+1}) = 0$ when $j_1 \neq j_3$. As a result, we have

$$\begin{aligned} \|\widehat{V} - V_0\|_{\infty, \infty} = & \max_{\substack{(j_1, j_2), (j_1, j_3) \\ \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})}} \left| \text{cov} \left[\eta_{j_1, j_2}^{*(\ell)}, \eta_{j_1, j_3}^{*(\ell)} | \{X_{i,j}\}_{1 \leq i \leq n, 0 \leq j \leq d+1} \right] \right. \\ & \left. - \text{cov} \left[\eta_{0, j_1, j_2}^{(\ell)}, \eta_{0, j_1, j_3}^{(\ell)} | \{X_{i,j}\}_{i \in \mathcal{I}_\ell^c, 0 \leq j \leq d+1} \right] \right|. \end{aligned}$$

After some calculations, we have, for $1 \leq j_1 \leq d+1$,

$$\begin{aligned} \text{cov} \left[\eta_{0, j_1, j_2}^{(\ell)}, \eta_{0, j_1, j_3}^{(\ell)} | \{X_{i,j}\}_{i \in \mathcal{I}_\ell^c, 0 \leq j \leq d+1} \right] &= \frac{\mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \boldsymbol{\beta}_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} | \widetilde{\mathbf{W}}^{(\ell)} \right] \sigma_*^2}{I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell)}, \\ \text{cov} \left[\eta_{j_1, j_2}^{*(\ell)}, \eta_{j_1, j_3}^{*(\ell)} | \{X_{i,j}\}_{1 \leq i \leq n, 0 \leq j \leq d+1} \right] &= \frac{|\mathcal{I}_\ell^c|^{-1} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left[\left\{ X_{i,k} - \widehat{\boldsymbol{\beta}}^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right] \widehat{\sigma}_*^2}{I_2(j_1, j_2, \ell) I_2(j_1, j_3, \ell)}. \end{aligned}$$

It follows that

$$\begin{aligned} \|\widehat{V} - V_0\|_{\infty, \infty} = & \max_{\substack{(j_1, j_2), (j_1, j_3) \\ \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})}} I_{3,1}(j_1, j_2, j_3, \ell) + \max_{\substack{(j_1, j_2), (j_1, j_3) \\ \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})}} I_{3,2}(j_1, j_2, j_3, \ell) \\ & + \max_{\substack{(j_1, j_2), (j_1, j_3) \\ \in \mathcal{S}(q_1, q_2, \widehat{\mathbf{B}}^{(\ell)})}} I_{3,3}(j_1, j_2, j_3, \ell), \end{aligned}$$

where

$$\begin{aligned}
I_{3,1}(j_1, j_2, j_3, \ell) &= \frac{\left| \mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \right|}{I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell)} |\widehat{\sigma}_*^2 - \sigma_*^2|, \\
I_{3,2}(j_1, j_2, j_3, \ell) &= \frac{\widehat{\sigma}_*^2}{I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell)} \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left[\left\{ X_{i,k} - \widehat{\beta}^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right] \right. \\
&\quad \left. - \mathbb{E} \prod_{k \in \{j_2, j_3\}} \left[\left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \right] \right|, \\
I_{3,3}(j_1, j_2, j_3, \ell) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left[\left\{ X_{i,k} - \widehat{\beta}^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right] \widehat{\sigma}_*^2 \right| \\
&\quad \times \frac{|I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell) - I_2(j_1, j_2, \ell) I_2(j_1, j_3, \ell)|}{I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell) I_2(j_1, j_2, \ell) I_2(j_1, j_3, \ell)}.
\end{aligned}$$

We next bound $\max_{j_1, j_2, j_3, \ell} I_{3,q}(j_1, j_2, j_3, \ell)$, $q = 1, 2, 3$, respectively.

For $\max_{j_1, j_2, j_3, \ell} I_{3,1}(j_1, j_2, j_3, \ell)$, by Cauchy-Schwarz inequality, we have

$$\begin{aligned}
&\max_{j_1, j_2, j_3, \ell} \left| \mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \right| \\
&\leq \max_{j_1, j_2, j_3, \ell} \prod_{k \in \{j_2, j_3\}} \mathbb{E} \left[\left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\}^2 \mid \widetilde{\mathbf{W}}^{(\ell)} \right].
\end{aligned}$$

By (48), we have,

$$\max_{j_1, j_2, j_3, \ell} \left| \mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \right| = O(1). \quad (50)$$

Combining (50) with (41) and Lemma 4 yields,

$$\max_{j_1, j_2, j_3, \ell} I_{3,1}(j_1, j_2, j_3, \ell) = O(n^{-\kappa_7}). \quad (51)$$

For $\max_{j_1, j_2, j_3, \ell} I_{3,2}(j_1, j_2, j_3, \ell)$, by Lemma 4, we have $\widehat{\sigma}_*^2 = O(1)$. Define

$$\begin{aligned} I_3^*(j_1, j_2, j_3, \ell) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left[\left\{ X_{i,k} - \widehat{\beta}^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right] \right. \\ &\quad \left. - \mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \right| \\ I_3^{**}(j_1, j_2, j_3) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_{i,k} - \widehat{\beta}^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right. \right. \\ &\quad \left. \left. - \prod_{k \in \{j_2, j_3\}} \left\{ X_{i,k} - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_i \right\} \right] \right|. \end{aligned}$$

Similar to (38), we have,

$$\begin{aligned} \max_{j_1, j_2, j_3, \ell} \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left\{ X_{i,k} - \beta_0^\top(j_1, k) \mathbf{X}_i \right\} - \mathbb{E} \left[\prod_{k \in \{j_2, j_3\}} \left\{ X_k - \beta_0^{(\ell)\top}(j_1, k) \mathbf{X}_0 \right\} \mid \widetilde{\mathbf{W}}^{(\ell)} \right] \right| \\ = O(n^{-1/2} \sqrt{\log n}), \end{aligned}$$

Next, $\max_{j_1, j_2, j_3, \ell} I_3^{**}(j_1, j_2, j_3, \ell) \leq I_{4,1}(j_1, j_2, j_3, \ell) + I_{4,2}(j_1, j_2, j_3, \ell) + I_{4,3}(j_1, j_2, j_3, \ell)$, where

$$\begin{aligned} I_{4,1}(j_1, j_2, j_3, \ell) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \beta_0^{(\ell)}(j_1, j_2) - \widehat{\beta}^{(\ell)}(j_1, j_2) \right\}^\top \mathbf{X}_i \left\{ X_{i,j_3} - \beta_0^{(\ell)\top}(j_1, j_3) \mathbf{X}_i \right\} \right|, \\ I_{4,2}(j_1, j_2, j_3, \ell) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \beta_0^{(\ell)}(j_1, j_3) - \widehat{\beta}^{(\ell)}(j_1, j_3) \right\}^\top \mathbf{X}_i \left\{ X_{i,j_2} - \beta_0^{(\ell)\top}(j_1, j_2) \mathbf{X}_i \right\} \right|, \\ I_{4,3}(j_1, j_2, j_3, \ell) &= \left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \left\{ \beta_0^{(\ell)}(j_1, j_2) - \widehat{\beta}^{(\ell)}(j_1, j_2) \right\}^\top \mathbf{X}_i \mathbf{X}_i^\top \left\{ \beta_0^{(\ell)}(j_1, j_2) - \widehat{\beta}^{(\ell)}(j_1, j_2) \right\} \right|. \end{aligned}$$

Using similar arguments in proving (35) and (40), we have

$$\begin{aligned} \max_{j_1, j_2, j_3, \ell} I_{4,1}(j_1, j_2, j_3, \ell) &= O(n^{-\kappa_5}), \\ \max_{j_1, j_2, j_3, \ell} I_{4,2}(j_1, j_2, j_3, \ell) &= O(n^{-\kappa_5}), \\ \max_{j_1, j_2, j_3, \ell} I_{4,3}(j_1, j_2, j_3, \ell) &= O(n^{-2\kappa_5}). \end{aligned}$$

Therefore, we have, for some constant $\kappa > 0$,

$$\max_{j_1, j_2, j_3, \ell} I_2^*(j_1, j_2, j_3, \ell) = O(n^{-\kappa}). \quad (52)$$

It then follows from (41) that, for some constant $\kappa > 0$,

$$\max_{j_1, j_2, j_3, \ell} I_2(j_1, j_2, j_3, \ell) = O(n^{-\kappa}). \quad (53)$$

For $\max_{j_1, j_2, j_3, \ell} I_{3,3}(j_1, j_2, j_3, \ell)$, combining (52) with (50), we have,

$$\left| \frac{1}{|\mathcal{I}_\ell^c|} \sum_{i \in \mathcal{I}_\ell^c} \prod_{k \in \{j_2, j_3\}} \left[\{X_{i,k} - \widehat{\beta}^{(\ell)\top}(j_1, k) \mathbf{X}_i\} \right] \right| = O(1). \quad (54)$$

Using similar arguments in bounding $\max_{j_1, j_2, \ell} |I_2(j_1, j_2, \ell) - I_2^{**}(j_1, j_2, \ell)|$ in Step 2, we get

$$\max_{j_1, j_2, j_3, \ell} |I_2^{**}(j_1, j_2, \ell) I_2^{**}(j_1, j_3, \ell) - I_2(j_1, j_2, \ell) I_2(j_1, j_3, \ell)| = O(n^{-\kappa}),$$

for some constant $\kappa > 0$. Combining this together with (54), (41), (43) and that $\widehat{\sigma}_*^2 = O(1)$, we have, for some constant $\kappa > 0$,

$$\max_{j_1, j_2, j_3, \ell} I_{3,3}(j_1, j_2, j_3, \ell) = O(n^{-\kappa}). \quad (55)$$

Combining (51), (53) and (55) together yields the bound for $\|\widehat{V} - V_0\|_{\infty, \infty}$ in (32). This completes Step 3. \square

.7 Proof of Theorem 2

Recall the proposed bootstrap procedure repeatedly generate random variables from $N(0, \widehat{V})$, and the critical value $\widehat{c}^{(\ell)}(q_1, q_2)$ is the upper $(\alpha/2)$ th quantile of $\|N(0, \widehat{V})\|_\infty$. That is,

$$\Pr \left(\|N(0, \widehat{V})\|_\infty \leq \widehat{c}^{(\ell)}(q_1, q_2) | \widehat{V} \right) = \frac{\alpha}{2}. \quad (56)$$

We have shown in the proof of Theorem 1 that the diagonal elements of \widehat{V} are uniformly bounded by some constant $\kappa > 0$. It follows from Bonferroni's inequality that

$$\begin{aligned} \Pr \left(\|N(0, \widehat{V})\|_\infty > t\sqrt{\kappa \log n} \right) &\leq (d+2) \max_{j \in \{0, \dots, d+1\}} \Pr \left(|N(0, \widehat{V}_{j,j})| > t\sqrt{\kappa \log n} \right) \\ &\leq (d+2) \left\{ 1 - \Phi \left(t\sqrt{\log n} \right) \right\}, \end{aligned}$$

where $\widehat{V}_{j,j}$ is the (j, j) th entry of \widehat{V} . For $t \geq 1$ and $n \geq 3$, we have $t\sqrt{\log n} \geq 1$, and hence

$$\begin{aligned} 1 - \Phi\left(t\sqrt{\log n}\right) &= \frac{1}{\sqrt{2\pi}} \int_{t\sqrt{\log n}}^{+\infty} \exp\left(-\frac{x^2}{2}\right) dx \leq \int_{t\sqrt{\log n}}^{+\infty} x \exp\left(-\frac{x^2}{2}\right) dx \\ &= \exp\left(-\frac{t^2 \log n}{2}\right) = n^{-t^2/2}. \end{aligned}$$

Setting $t = 2\sqrt{\kappa_1 + 1}$, it follows from the condition $d = O(n^{\kappa_1})$ that,

$$\Pr\left(\|N(0, \widehat{V})\|_{\infty} > \sqrt{2(\kappa_1 + 1)c \log n}\right) = O(d/n^{\kappa_1+1}) = o(1).$$

In view of (56), we obtain $\widehat{c}^{(\ell)}(q_1, q_2) \leq \sqrt{2(\kappa_1 + 1)c \log n}$.

According to our test procedure, we reject the null if $\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0} > \widehat{c}(0, q)$ and $\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{d+1,q} > \widehat{c}(q, d+1)$, for some $\ell = 1, 2$, where $\widehat{\mathbf{W}}^{*(\ell)} = |\widehat{\mathbf{W}}^{(\ell)}| \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(2)} \oplus \dots \oplus |\widehat{\mathbf{W}}^{(\ell)}|^{(d)}$. To prove Theorem 2, it suffices to show

$$\begin{aligned} \Pr(\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0} \geq \sqrt{2(\kappa_1 + 1)c \log n}) &\rightarrow 1, \\ \Pr(\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{d+1,q} \geq \sqrt{2(\kappa_1 + 1)c \log n}) &\rightarrow 1, \end{aligned} \tag{57}$$

under the alternative hypothesis H_1 . We next prove (57).

Under the given conditions of the theorem, there exists a path, $X_0 \rightarrow X_{i_1} \rightarrow \dots \rightarrow X_{i_k} \rightarrow X_q$, such that $|W_{0,0,i_1}|, |W_{0,i_1,i_2}|, \dots, |W_{0,i_{k-1},i_k}|, |W_{0,i_k,q}| \gg n^{-1/2}\sqrt{\log n}$. Under the given conditions, we have with probability tending to 1 that,

$$\max\left(|\widehat{W}_{0,i_1}^{(\ell)} - W_{0,0,i_1}|, \max_{1 \leq j \leq k-1} |\widehat{W}_{i_j,i_{j+1}}^{(\ell)} - W_{0,i_j,i_{j+1}}|, |\widehat{W}_{i_k,q}^{(\ell)} - W_{0,i_k,q}|\right) = O(n^{-1/2}\sqrt{\log n}).$$

It follows that

$$\min\left(|\widehat{W}_{0,i_1}^{(\ell)}|, \min_{1 \leq j \leq k-1} |\widehat{W}_{i_j,i_{j+1}}^{(\ell)}|, |\widehat{W}_{i_k,q}^{(\ell)}|\right) \gg n^{-1/2}\sqrt{\log n}. \tag{58}$$

By definition, we have that,

$$\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0} \geq \sqrt{n} \min\left(|\widehat{W}_{0,i_1}^{(\ell)}|, \min_{1 \leq j \leq k-1} |\widehat{W}_{i_j,i_{j+1}}^{(\ell)}|, |\widehat{W}_{i_k,q}^{(\ell)}|\right).$$

It follows from (58) that $\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0} \gg \sqrt{\log n}$ under Condition (A5). Therefore, $\Pr\left(\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0} \geq \sqrt{2(\kappa_1 + 1)c \log n}\right) \rightarrow 1$, which proves the first result in (57). Similarly, we can prove the second result in (57) that $\Pr\left(\sqrt{n}(\widehat{\mathbf{W}}^{*(\ell)})_{d+1,q} \geq \sqrt{2(\kappa_1 + 1)c \log n}\right) \rightarrow 1$. This completes the proof of Theorem

2.

□

.8 Proof of Theorem 3

First, we show that the p -values $\widehat{p}^{(\ell)}(0, q)$ and $\widehat{p}^{(\ell)}(q, d+1)$ are asymptotically independent. Let $\text{ACT}(q)$ and $\text{DES}(q)$ denote the set of ancestors and descendants of X_q , respectively. By Condition (A4), the two test statistics $\sqrt{\mathcal{I}_\ell^c}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0}$ and $\sqrt{\mathcal{I}_\ell^c}(\widehat{\mathbf{W}}^{*(\ell)})_{d+1,q}$ are constructed based on the set of estimators $\left\{\widehat{W}_{j_1,j_2}^{(\ell)} : j_1 \in \text{ACT}(q) \cup \{q\}, j_2 \in \text{ACT}(q)\right\}$ and $\left\{\widehat{W}_{j_1,j_2}^{(\ell)} : j_1 \in \text{DES}(q), j_2 \in \text{DES}(q) \cup \{q\}\right\}$, respectively. These two sets of estimators are asymptotically independent conditional on $\{X_i\}_{i \in \mathcal{I}_\ell}$. To better illustrate this, note that for any (j_1, j_2, j_3, j_4) such that $j_1 \in \text{ACT}(q) \cup \{q\}, j_2 \in \text{ACT}(q), j_3 \in \text{DES}(q), j_4 \in \text{DES}(q) \cup \{q\}$, we have $j_1 \neq j_3$. As discussed in Section 4.4, $\sqrt{n}(\widehat{W}_{j_1,j_2}^{(\ell)} - W_{0,j_1,j_2})$ and $\sqrt{n}(\widehat{W}_{j_3,j_4}^{(\ell)} - W_{0,j_3,j_4})$ are asymptotically uncorrelated. Since the two variables are jointly normal, they are asymptotically independent as well. As a result, the two test statistics $\sqrt{\mathcal{I}_\ell^c}(\widehat{\mathbf{W}}^{*(\ell)})_{q,0}$ and $\sqrt{\mathcal{I}_\ell^c}(\widehat{\mathbf{W}}^{*(\ell)})_{d+1,q}$ are asymptotically independent, and so are their corresponding p -values.

Next, a closer look at the proof of Theorem 1 shows that the type-I error rates can be uniformly controlled across different mediators. That is,

$$\max_{\ell \in \{1,2\}} \max_{q \in \{1,\dots,d\}} \Pr\left\{\widehat{p}_{\max}^{(\ell)}(q) \leq \alpha \mid H_0(q) \text{ holds}\right\} \leq \alpha + o(1), \quad (59)$$

for any given significance level $0 < \alpha < 1$. Following the proof of Theorem 1 in Djordjilović et al. (2019), and by (59), we can show that,

$$\max_{\ell \in \{1,2\}} \max_{q \in \{1,\dots,d\}} \Pr\left\{\widehat{p}_{\max}^{(\ell)}(q) \leq \alpha \mid H_0(q) \text{ holds}, \widehat{p}_{\min}^{(\ell)}(q) \leq c^{(\ell)}\right\} \leq \alpha + o(1), \quad (60)$$

for any significance level α and the critical value $c^{(\ell)} > 0$.

Then, similar to Theorem 1.3 of Benjamini and Yekutieli (2001b), and by (60), we have that $\text{FDR}(\mathcal{H}^{(\ell)})$ is guaranteed at level $\alpha/2$. Recall $\mathcal{H} = \mathcal{H}^{(1)} \cup \mathcal{H}^{(2)}$ is the final set of our selected mediators, and \mathcal{N} is the set of unimportant mediators. It follows that,

$$\begin{aligned} \text{FDR}(\mathcal{H}) &= \mathbb{E}\left(\frac{|\mathcal{N} \cup \mathcal{H}|}{\max(1, \mathcal{H})}\right) \leq \sum_{\ell \in \{1,2\}} \mathbb{E}\left(\frac{|\mathcal{N} \cup \mathcal{H}^{(\ell)}|}{\max(1, \mathcal{H})}\right) \leq \sum_{\ell \in \{1,2\}} \mathbb{E}\left(\frac{|\mathcal{N} \cup \mathcal{H}^{(\ell)}|}{\max(1, \mathcal{H}^{(\ell)})}\right) \\ &= \sum_{\ell \in \{1,2\}} \text{FDR}(\mathcal{H}^{(\ell)}) = \alpha + o(1). \end{aligned}$$

This completes the proof of Theorem 3. □

.9 Proof of Proposition 1

Outline of the proof: We choose $\lambda = \kappa\sqrt{\log n/n}$ for some sufficiently large constant $\kappa > 0$. We divide the proof into three steps. In Step 1, we establish the uniform convergence rate of $\|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi) - \mathbf{W}_j(\pi)\|_2$. Due to the acyclic constraint, $\widetilde{\mathbf{W}}^{(\ell)} = \widetilde{\mathbf{W}}^{(\ell)}(\pi)$ for some ordering π . In Step 2, we show that

$$\begin{aligned} & \min_{\pi \notin \Pi_0} \left(\sum_{i \in \mathcal{I}_\ell} \|\mathbf{X}_i - \widetilde{\mathbf{W}}^{(\ell)}(\pi) \mathbf{X}_i\|^2 + \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 \right) \\ & > \max_{\pi \in \Pi_0} \left(\sum_{i \in \mathcal{I}_\ell} \|\mathbf{X}_i - \mathbf{W}(\pi) \mathbf{X}_i\|^2 + \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\mathbf{W}_j(\pi)\|_1 \right). \end{aligned} \quad (61)$$

By definition, this implies that $\widetilde{\mathbf{W}}^{(\ell)} = \widetilde{\mathbf{W}}^{(\ell)}(\pi^*)$ for some true ordering $\pi^* \in \Pi^*$. This proves the first assertion of Proposition 1. In the third step, we derive the convergence rate of $\widetilde{\mathbf{W}}^{(\ell)}$, and completes the second assertion of Proposition 1.

Step 1: For any π, i, j , let $\varepsilon_{i,j}(\pi) = X_{i,j} - \mathbf{X}_i^\top \mathbf{W}_j(\pi)$ denote the residual. By definition,

$$\begin{aligned} & \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi)|^2 + \lambda |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 \\ & \leq \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |X_{i,j} - \mathbf{X}_i^\top \mathbf{W}_j(\pi)|^2 + \lambda |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \|\mathbf{W}_j(\pi)\|_1 = \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + \lambda |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \|\mathbf{W}_j(\pi)\|_1. \end{aligned} \quad (62)$$

Note that

$$\begin{aligned} & \sum_{i \in \mathcal{I}_\ell} |X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi)|^2 = \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi) + \mathbf{X}_i^\top \mathbf{W}_j(\pi) - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi)|^2 \\ & = \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + 2 \sum_{i \in \mathcal{I}_\ell} \varepsilon_{i,j}(\pi) \left\{ \mathbf{X}_i^\top \mathbf{W}_j(\pi) - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} + \sum_{i \in \mathcal{I}_\ell} \left[\mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2. \end{aligned}$$

Let $s_j(\pi)$ and $\tilde{s}_j^{(\ell)}(\pi)$ denote the number of nonzero elements in $\mathbf{W}_j(\pi)$ and $\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)$, respectively. By Theorem 7.1 of van de Geer and Bühlmann (2013), we have,

$$\begin{aligned} & \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} 2 \left| \varepsilon_{i,j}(\pi) \mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right| - \delta \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} \left| \mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right|^2 \\ & \leq \sum_{j=0}^{d+1} \frac{\kappa \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n}{\delta}, \end{aligned} \quad (63)$$

for some constant $\kappa > 0$, any $\delta > 0$, and any order π . Under the event defined in (63),

$$\begin{aligned} \sum_{i \in \mathcal{I}_\ell} \left| X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right|^2 &\geq \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 \\ &+ (1 - \delta) \sum_{i \in \mathcal{I}_\ell} \sum_{j=0}^{d+1} \left[\mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2 - \sum_{j=0}^{d+1} \frac{\kappa \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n}{\delta}. \end{aligned} \quad (64)$$

Setting $\delta = 1/2$, together with (62), we have,

$$\begin{aligned} &\frac{1}{2} \sum_{i \in \mathcal{I}_\ell} \sum_{j=0}^{d+1} \left[\mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2 + \lambda |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \left\| \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\|_1 \\ &\leq \sum_{j=0}^{d+1} 2\kappa \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n + \lambda |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \left\| \mathbf{W}_j(\pi) \right\|_1. \end{aligned}$$

Using similar arguments in bounding $\max_{j_1, j_2, \ell} I_{N,1}(j_1, j_2, \ell)$ in Step 1 of the proof of Theorem 1, we can show there exists constants $\kappa^* \leq \lambda_{\min}(\Sigma_0) \leq \lambda_{\max}(\Sigma_0) \leq 1/\kappa^*$ such that

$$\begin{aligned} \kappa^* |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \left\| \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\|_2^2 &\leq \frac{1}{2} \sum_{i \in \mathcal{I}_\ell} \sum_{j=0}^{d+1} \left[\mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2 \\ &\leq \frac{1}{\kappa^*} |\mathcal{I}_\ell| \sum_{j=0}^{d+1} \left\| \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\|_2^2, \end{aligned} \quad (65)$$

for any order π . It then follows that, for any order π ,

$$\begin{aligned} \kappa^* \sum_{j=0}^{d+1} \left\| \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\|_2^2 + \lambda \sum_{j=0}^{d+1} \left\| \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\|_1 &\leq \sum_{j=0}^{d+1} \frac{4c}{n} \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n \\ &+ \lambda \sum_{j=0}^{d+1} \left\| \mathbf{W}_j(\pi) \right\|_1. \end{aligned} \quad (66)$$

Next, define the $(d+2)$ -dimensional vector,

$$\boldsymbol{\mu}(\pi) = \left\{ \widetilde{\mathbf{W}}_0^{(\ell)\top}(\pi) - \mathbf{W}_0^{(\ell)\top}(\pi), \widetilde{\mathbf{W}}_1^{(\ell)\top}(\pi) - \mathbf{W}_1^{(\ell)\top}(\pi), \dots, \widetilde{\mathbf{W}}_{d+1}^{(\ell)\top}(\pi) - \mathbf{W}_{d+1}^{(\ell)\top}(\pi) \right\}^\top.$$

Let $\mathcal{M}(\pi)$ denote the support of $\boldsymbol{\mu}(\pi)$. Our goal is to bound $\|\boldsymbol{\mu}(\pi)\|_2 = \sum_{j=0}^{d+1} \left\| \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) - \mathbf{W}_j(\pi) \right\|_2^2$. We next consider two cases.

First, when $\|\boldsymbol{\mu}(\pi)\|_2 \leq \sum_{j=0}^{d+1} 4c \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n / (n\kappa^*)$, it is already bounded.

Second, when $\|\boldsymbol{\mu}(\pi)\|_2 > \sum_{j=0}^{d+1} 4c \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n / (n\kappa^*)$, we aim to obtain a

larger upper bound. Specifically, under the event defined in (66), we have, for all π ,

$$\sum_{j=0}^{d+1} \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 \leq \sum_{j=0}^{d+1} \|\mathbf{W}_j(\pi)\|_1.$$

It follows that, for all π ,

$$\|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_1 \leq \|\boldsymbol{\mu}(\pi)_{\mathcal{M}^c(\pi)}\|_1, \quad (67)$$

where $\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}$ and $\boldsymbol{\mu}(\pi)_{\mathcal{M}^c(\pi)}$ denote the sub-vector formed by the elements in $\mathcal{M}(\pi)$ and $\mathcal{M}^c(\pi)$, respectively. Then, under the event defined in (67), and by the definition of s_0 ,

$$\begin{aligned} \sum_{j=0}^{d+1} \|\mathbf{W}_j(\pi)\|_1 - \sum_{j=0}^{d+1} \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 &\leq \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_1 + \|\boldsymbol{\mu}(\pi)_{\mathcal{M}^c(\pi)}\|_1 \leq 2\|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_1 \\ &\leq 2\sqrt{|\mathcal{M}(\pi)|} \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2 \leq 2\sqrt{s_0(d+2)} \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2. \end{aligned}$$

By (66), we have that, for all π ,

$$\begin{aligned} \kappa^* \|\boldsymbol{\mu}(\pi)\|_2^2 &\leq \sum_{j=0}^{d+1} \frac{4c}{n} \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n + 2\lambda\sqrt{s_0(d+2)} \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2, \\ \kappa^* \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2^2 &\leq \sum_{j=0}^{d+1} \frac{4c}{n} \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n + 2\lambda\sqrt{s_0(d+2)} \|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2. \end{aligned}$$

It then follows that, for all π and some positive constant $O(1)$,

$$\|\boldsymbol{\mu}(\pi)_{\mathcal{M}(\pi)}\|_2 \leq O(1) \left(\lambda\sqrt{s_0(d+2)} + \sqrt{\sum_{j=0}^{d+1} \frac{4c}{n} \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n} \right).$$

Under the event defined in (67), and by Equations (A5), (A6) in Zhou (2009), we have,

$$\|\boldsymbol{\mu}(\pi)\|_2 \leq O(1) \left(\lambda\sqrt{s_0(d+2)} + \sqrt{\sum_{j=0}^{d+1} \frac{4c}{n} \left\{ s_j(\pi) + \tilde{s}_j^{(\ell)}(\pi) \right\} \log n} \right), \quad (68)$$

for all π .

Therefore, in both cases, we show that $\|\boldsymbol{\mu}(\pi)\|_2$ is uniformly bounded for all π . This completes Step 1.

Step 2: Under the events defined in (65) and (68), we have that, for all π ,

$$\sum_{i \in \mathcal{I}_\ell} \sum_{j=0}^{d+1} \left[\mathbf{x}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2 \leq O(1) d^2 \log n.$$

Combining this together with (64) yields that, for all π ,

$$\sum_{i \in \mathcal{I}_\ell} \left| X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right|^2 \geq \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + O(1)d^2 \log n. \quad (69)$$

Under the omega-min condition in (A6), using similar arguments in the proof of Lemma 7.8 of van de Geer and Bühlmann (2013), we have,

$$\min_{\pi \notin \Pi_0} \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 \geq \max_{\pi \in \Pi_0} \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + \kappa(d+2)n,$$

for some constant $\kappa > 0$. Under the event defined in (69), we have,

$$\begin{aligned} \min_{\pi \notin \Pi_0} \left\{ \sum_{i \in \mathcal{I}_\ell} \|\mathbf{X}_i - \widetilde{\mathbf{W}}^{(\ell)}(\pi) \mathbf{X}_i\|^2 + \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 \right\} &\geq \min_{\pi \notin \Pi_0} \sum_{i \in \mathcal{I}_\ell} \|\mathbf{X}_i - \widetilde{\mathbf{W}}^{(\ell)}(\pi) \mathbf{X}_i\|^2 \\ &\geq \max_{\pi \in \Pi_0} \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + \kappa(d+2)n - O(1)d^2 \log n. \end{aligned}$$

Under the Condition $\|\mathbf{W}_0\|_2 = O(1)$ for any $\pi \in \Pi^*$, we have $\|\mathbf{W}_j(\pi)\|_2 = \|\mathbf{W}_{0,j}\|_2 = O(1)$.

As a result,

$$\max_{\pi \in \Pi_0} \sum_{j=0}^{d+1} \lambda |\mathcal{I}_\ell| \|\mathbf{W}_j(\pi)\|_1 = \sum_{j=0}^{d+1} \lambda |\mathcal{I}_\ell| \|\mathbf{W}_{0,j}\|_1 \leq \sqrt{d+2} \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\mathbf{W}_{0,j}\|_2 = O\left(d^{3/2} n^{1/2} \sqrt{\log n}\right).$$

It follows that,

$$\begin{aligned} &\min_{\pi \notin \Pi_0} \left(\sum_{i \in \mathcal{I}_\ell} \|\mathbf{X}_i - \widetilde{\mathbf{W}}^{(\ell)}(\pi) \mathbf{X}_i\|^2 + \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\widetilde{\mathbf{W}}_j^{(\ell)}(\pi)\|_1 \right) \\ &\geq \max_{\pi \in \Pi_0} \sum_{j=0}^{d+1} \left(\sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + \sum_{j=0}^{d+1} |\mathcal{I}_\ell| \lambda \|\mathbf{W}_j(\pi)\|_1 \right) \\ &\quad + \kappa(d+2)n - O(1)d^2 \log n - O\left(d^{3/2} n^{1/2} \sqrt{\log n}\right), \end{aligned}$$

where the last line is strictly positive. As a result, we obtain (61). This completes Step 2.

Step 3: Built on the results obtained in Step 2, we have $\hat{\pi} \in \Pi_0$, where $\hat{\pi}$ is the estimated ordering. To complete the proof, it suffices to show that $\max_{\pi \in \Pi_0} \max_{j \in \{0, \dots, d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2 = O\left(n^{-1/2} \sqrt{s_0 \log n}\right)$. Since the number of elements in Π^* is bounded by $O(n^{\kappa^*})$ for

some $\kappa^* > 0$, by Bernstein's inequality and Bonferroni's inequality, we have that,

$$\max_{\pi \in \Pi_0} \max_{k \in \{0,1,\dots,j-1\}} \left| \sum_{i \in \mathcal{I}_\ell} \varepsilon_{i,j}(\pi) X_{i,\pi_k} \right| = O\left(n^{1/2} \sqrt{\log n}\right).$$

The explicit convergence rate of $\max_{\pi \in \Pi_0} \max_{j \in \{0,\dots,d+1\}} \|\widetilde{\mathbf{W}}_j^{(\ell)} - \mathbf{W}_{0,j}\|_2$ can then be similarly derived following Theorem 3.1 of Zhou (2009). This completes Step 3.

Remark: Finally, we discuss the possibility of relaxing the condition $d = o(n)$, by imposing some sparsity conditions on $\widetilde{\mathbf{W}}^{(\ell)}(\pi)$, and the population limit of $\widetilde{\mathbf{W}}^{(\ell)}(\pi)$, i.e., $\mathbf{W}_0(\pi)$. In this case, the condition $d \ll n$ can be relaxed to $d = O(n^{\kappa_1})$, for some constant $\kappa_1 > 0$. Since we do not require $\kappa_1 < 1$, the dimension d can grow at a much faster rate than n .

Specifically, suppose the number of nonzero elements in each column of $\widetilde{\mathbf{W}}^{(\ell)}(\pi)$ and $\mathbf{W}_0(\pi)$ satisfies that $\max_{j,\pi} [\max\{s_j(\pi), \tilde{s}_j^{(\ell)}(\pi)\}] = O(n^{\kappa_*})$, for some constant $0 < \kappa_* < 1$. It follows from (64) and (68) that

$$\begin{aligned} \sum_{i \in \mathcal{I}_\ell} \left| X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right|^2 &\geq \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + (1 - \delta) \sum_{i \in \mathcal{I}_\ell} \sum_{j=0}^{d+1} \left[\mathbf{X}_i^\top \left\{ \mathbf{W}_j(\pi) - \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right\} \right]^2 \\ &\quad - \sum_{j=0}^{d+1} \frac{O(1)n^{\kappa_*} \log n}{\delta}, \\ \|\boldsymbol{\mu}(\pi)\|_2 &\leq O(1) \left(\lambda \sqrt{dn^{\kappa_*/2}} + \sqrt{dn^{\kappa_*/2-1/2}} \sqrt{\log n} \right), \end{aligned}$$

for all π , where $O(1)$ denotes some positive constant. Therefore, we have

$$\sum_{i \in \mathcal{I}_\ell} \left| X_{i,j} - \mathbf{X}_i^\top \widetilde{\mathbf{W}}_j^{(\ell)}(\pi) \right|^2 \geq \sum_{j=0}^{d+1} \sum_{i \in \mathcal{I}_\ell} |\varepsilon_{i,j}(\pi)|^2 + O(1)dn^{\kappa_*} \log n,$$

for all π . The rest of the proof follows using similar arguments as in Steps 2 and 3. \square

.10 Additional numerical results

First, we provide more details about the simulation setup. Table 2 reports the corresponding mediators with nonzero mediation effects, and their associated $\delta(q)$, where $\delta(q) = (\mathbf{W}_0^*)_{d+1,q}(\mathbf{W}_0^*)_{q,0}$, \mathbf{W}_0^* is constructed based on \mathbf{W}_0 , $q = 1, \dots, d$. Meanwhile, Figure 6 shows an example of the adjacent matrix when $d = 100$.

Next, we report the simulation results for testing of a single mediator. Figures 2 and 8 show the results for $d = 100$ and $d = 150$, respectively, which complement the results for

Scenario A: $(d, p_1, p_2) = (50, 0.05, 0.15)$, $n = 100, 200$										
q	10	12	20	28	30	41				
$\delta(q)$	1.06	1.03	0.63	1.08	0.64	1.31				
Scenario B: $(d, p_1, p_2) = (100, 0.025, 0.075)$, $n = 250, 500$										
q	5	14	17	20	30	71	80	82	89	
$\delta(q)$	0.99	0.33	0.92	0.85	0.85	1.37	0.72	0.72	0.72	
Scenario C: $(d, p_1, p_2) = (150, 0.02, 0.05)$, $n = 250, 500$										
q	19	23	51	52	54	56	60	81	92	93
$\delta(q)$	1.24	0.93	1.71	1.28	1.90	1.90	0.94	0.88	0.88	1.02

Table 2: Mediators with nonzero mediation effects, $q = 1, \dots, d$.

$d = 50$ reported in Section 6 of the paper. We observe a similar qualitative pattern that our proposed test achieves a valid size under the null hypothesis, and achieves a larger empirical power under the alternative hypothesis.

Next, we report the simulation results for multiple testing. Figure 9 show the results for $d = 100$ and $d = 150$, respectively, which complement the results for $d = 50$ reported in Section 6 of the paper. We again observe a similar qualitative pattern that both our test and the standard Benjamini-Yekutieli (BY) procedure achieve a valid false discovery control, whereas our method is more powerful than BY.

Finally, we evaluate the constant variance condition for the real data example. Following a similar idea of Li et al. (2019), we compute and plot the residuals $X_{i,j} - \widetilde{\mathbf{W}}^{(\ell)} \mathbf{X}_i$, for $i \in \mathcal{I}_\ell^c$, $j \in \{0, \dots, d+1\}$, $\ell = 1, 2$. Figure 10 shows the boxplots of such residuals for the amyloid




Figure 6: The example adjacent matrix with $d = 100$.




Figure 7: Empirical rejection rate and ROC curve of the proposed test, LOGAN, and the test of Chakraborty et al. (2018), MIDA, when $d = 100$. The upper panels: $n = 250$, and the bottom panels: $n = 500$. The left panels: under H_0 , the middle panels: under H_1 , where the horizontal axis is the mediator index, and the right panels: the average ROC curve.




Figure 8: Empirical rejection rate and ROC curve of the proposed test, LOGAN, and the test of Chakraborty et al. (2018), MIDA, when $d = 150$. The upper panels: $n = 250$, and the bottom panels: $n = 500$. The left panels: under H_0 , the middle panels: under H_1 , where the horizontal axis is the mediator index, and the right panels: the average ROC curve.




Figure 9: False discover rate and true positive rate of the proposed method and the Benjamini-Yekutieli procedure. The horizontal axis corresponds to the significance level α . The upper panels: $d = 100$, and the bottom panels: $d = 150$. The left two panels: $n = 250$, and the right two panels: $n = 500$.

negative and amyloid positive groups, respectively. It is seen that the lengths of boxes and whiskers are similar across different variables, suggesting no strong evidence against the constant variance condition.




Figure 10: Residual plots for model diagnosis. The top panel: the amyloid negative group, and the bottom panel: the amyloid positive group.

References

- Alzheimer’s Association (2020). 2020 Alzheimer’s disease facts and figures. *Alzheimer’s & Dementia*, 16(3):391–460.
- Bakkour, A., Morris, J. C., Wolk, D. A., and Dickerson, B. C. (2013). The effects of aging and alzheimer’s disease on cerebral cortical anatomy: Specificity and differential relationships with cognition. *NeuroImage*, 76:332–344.
- Barber, R. F. and Candés, E. J. (2019). A knockoff filter for high-dimensional selective inference. *Ann. Statist.*, 47(5):2504–2537.
- Baron, R. M. and Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6):1173–1182.
- Benjamini, Y. and Yekutieli, D. (2001a). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pages 1165–1188.
- Benjamini, Y. and Yekutieli, D. (2001b). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165–1188.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732.
- Boca, S. M., Sinha, R., Cross, A. J., Moore, S. C., and Sampson, J. N. (2014). Testing multiple biological mediators simultaneously. *Bioinformatics*, 30(2):214–220.
- Candes, E. and Tao, T. (2007). The dantzig selector: Statistical estimation when p is much larger than n . *Annals of Statistics*, 35(6):2313–2351.
- Chakraborty, A., Nandy, P., and Li, H. (2018). Inference for individual mediation effects and interventional effects in sparse high-dimensional causal graphical models. *arXiv preprint arXiv:1809.10652*.
- Chakradhar, S. T., Agrawal, V. D., and Rothweiler, S. G. (1993). A transitive closure algorithm for test generation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 12(7):1015–1028.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters: Double/debiased machine learning. *The Econometrics Journal*, 21:C1–C68.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2013a). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2013b). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819.

- Chernozhukov, V., Chetverikov, D., and Kato, K. (2014). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics*, 42(4):1564–1597.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2015). Comparison and anti-concentration bounds for maxima of Gaussian random vectors. *Probability Theory and Related Fields*, 162(1-2):47–70.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2017). Detailed proof of nazarov’s inequality. *arXiv preprint arXiv:1711.10696*.
- Chickering, D. M. (2003). Optimal structure identification with greedy search. *Journal of Machine Learning Research (JMLR)*, 3(3):507–554. Computational learning theory.
- Djordjilović, V., Hemerik, J., and Thoresen, M. (2019). Optimal two-stage testing of multiple mediators. *arXiv preprint arXiv:1911.00862*.
- Djordjilović, V., Page, C. M., Gran, J. M., Nøst, T. H., Sandanger, T. M., Veierød, M. B., and Thoresen, M. (2019). Global test for high-dimensional mediation: testing groups of potential mediators. *Statistics in Medicine*, 38(18):3346–3360.
- Donohue, M. C., Sperling, R. A., and Others (2014). The Preclinical Alzheimer Cognitive Composite: Measuring Amyloid-Related Decline. *Journal of the American Medical Association: Neurology*, 71(8):961–970.
- Du, A.-T., Schuff, N., Kramer, J. H., Rosen, H. J., Gorno-Tempini, M. L., Rankin, K., Miller, B. L., and Weiner, M. W. (2007). Different regional patterns of cortical thinning in Alzheimer’s disease and frontotemporal dementia. *Brain*, 130(4):1159–1166.
- 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. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society. Series B.*, 70(5):849–911.
- Huang, Y.-T. (2018). Joint significance tests for mediation effects of socioeconomic adversity on adiposity via epigenetics. *The Annals of Applied Statistics*, 12(3):1535–1557.
- Huang, Y.-T. and Pan, W.-C. (2016). Hypothesis test of mediation effect in causal mediation model with high-dimensional continuous mediators. *Biometrics*, 72(2):401–413.
- Landau, S., Lu, M., Joshi, A., Pontecorvo, M., Mintun, M., Trojanowski, J., Shaw, L., and Jagust, W. (2013). Comparing pet imaging and csf measurements of a. *Annals of neurology*, 74.
- Li, C., Shen, X., and Pan, W. (2019). Likelihood ratio tests for a large directed acyclic graph. *Journal of the American Statistical Association*, 0(0):1–16.
- MacKinnon, D. P. and Fairchild, A. J. (2009). Current directions in mediation analysis. *Current Directions in Psychological Science*, 18(1):16–20.

- Meinshausen, N., Meier, L., and Bühlmann, P. (2009). p -values for high-dimensional regression. *Journal of the American Statistical Association*, 104(488):1671–1681.
- Nandy, P., Maathuis, M. H., and Richardson, T. S. (2017). Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. *Annals of Statistics*, 45(2):647–674.
- Newey, W. K. and Robins, J. R. (2018). Cross-fitting and fast remainder rates for semi-parametric estimation. *arXiv preprint arXiv:1801.09138*.
- Ning, Y. and Liu, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *The Annals of Statistics*, 45(1):158–195.
- Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 411–420.
- Peters, J. and Bühlmann, P. (2014). Identifiability of Gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228.
- Romano, J. and DiCiccio, C. (2019). Multiple data splitting for testing. Technical report, Stanford University Technical Report.
- Rubin, D., Dudoit, S., and van der Laan, M. (2006). A method to increase the power of multiple testing procedures through sample splitting. *Statistical applications in genetics and molecular biology*, 5:1544–6115.1148.
- Sampson, J. N., Boca, S. M., Moore, S. C., and Heller, R. (2018). Fwer and fdr control when testing multiple mediators. *Bioinformatics*, 34(14):2418–2424.
- Shi, C., Fan, A., Song, R., and Lu, W. (2018). High-dimensional A -learning for optimal dynamic treatment regimes. *The Annals of Statistics*, 46(3):925–957.
- Shi, C., Rui, S., Lu, W., and Li, R. (2020). Statistical inference for high-dimensional models via recursive online-score estimation. *Journal of the American Statistical Association*, accepted.
- Shi, C., Song, R., Chen, Z., Li, R., et al. (2019). Linear hypothesis testing for high dimensional generalized linear models. *The Annals of Statistics*, 47(5):2671–2703.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society. Series B.*, 73(3):273–282.
- van de Geer, S. and Bühlmann, P. (2013). ℓ_0 -penalized maximum likelihood for sparse directed acyclic graphs. *The Annals of Statistics*, 41(2):536–567.
- van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, 42(3):1166–1202.

- van Hoesen, G. W., Hyman, B. T., and Damasio, A. R. (1991). Entorhinal cortex pathology in alzheimer’s disease. *Hippocampus*, 1(1):1–8.
- Yuan, Y., Shen, X., Pan, W., and Wang, Z. (2019). Constrained likelihood for reconstructing a directed acyclic Gaussian graph. *Biometrika*, 106(1):109–125.
- Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942.
- Zhang, C.-H. and Zhang, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society. Series B.*, 76(1):217–242.
- Zhang, H., Zheng, Y., Zhang, Z., Gao, T., Joyce, B., Yoon, G., Zhang, W., Schwartz, J., Just, A., Colicino, E., et al. (2016). Estimating and testing high-dimensional mediation effects in epigenetic studies. *Bioinformatics*, 32(20):3150–3154.
- Zhao, Y. and Luo, X. (2016). Pathway lasso: Estimate and select sparse mediation pathways with high dimensional mediators. *arXiv preprint arXiv:1603.07749*.
- Zheng, X., Aragam, B., Ravikumar, P. K., and Xing, E. P. (2018). Dags with no tears: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, pages 9472–9483.
- Zhong, K., Yen, I. E.-H., Dhillon, I. S., and Ravikumar, P. K. (2014). Proximal quasi-newton for computationally intensive l1-regularized m-estimators. In *Advances in Neural Information Processing Systems*, pages 2375–2383.
- Zhou, S. (2009). Restricted eigenvalue conditions on subgaussian random matrices. *arXiv:0912.4045*.
- Zhu, Y., Shen, X., and Pan, W. (2020). On high-dimensional constrained maximum likelihood inference. *Journal of the American Statistical Association*, 115(529):217–230.