

Testing Conditional Independence via Quantile Regression Based Partial Copulas

Lasse Petersen

*Department of Mathematical Sciences
University of Copenhagen
Universitetsparken 5, 2100 Copenhagen, Denmark*

LP@MATH.KU.DK

Niels Richard Hansen

*Department of Mathematical Sciences
University of Copenhagen
Universitetsparken 5, 2100 Copenhagen, Denmark*

NIELS.R.HANSEN@MATH.KU.DK

Editor: Peter Spirtes

Abstract

The partial copula provides a method for describing the dependence between two random variables X and Y conditional on a third random vector Z in terms of nonparametric residuals U_1 and U_2 . This paper develops a nonparametric test for conditional independence by combining the partial copula with a quantile regression based method for estimating the nonparametric residuals. We consider a test statistic based on generalized correlation between U_1 and U_2 and derive its large sample properties under consistency assumptions on the quantile regression procedure. We demonstrate through a simulation study that the resulting test is sound under complicated data generating distributions. Moreover, in the examples considered the test is competitive to other state-of-the-art conditional independence tests in terms of level and power, and it has superior power in cases with conditional variance heterogeneity of X and Y given Z .

Keywords: Conditional independence testing, nonparametric testing, partial copula, conditional distribution function, quantile regression

1. Introduction

This paper introduces a new class of nonparametric tests of conditional independence between real-valued random variables, $X \perp\!\!\!\perp Y \mid Z$, based on quantile regression. Conditional independence is an important concept in many statistical fields such as graphical models and causal inference (Lauritzen, 1996; Spirtes et al., 2000; Pearl, 2009). However, Shah and Peters (2020) proved that conditional independence is an untestable hypothesis when the distribution of (X, Y, Z) is only assumed to be absolutely continuous with respect to Lebesgue measure.

More precisely, let \mathcal{P} denote the set of distributions of (X, Y, Z) that are absolutely continuous with respect to Lebesgue measure. Let $\mathcal{H} \subset \mathcal{P}$ be those distributions for which conditional independence holds. Then Shah and Peters (2020) showed that if ψ_n is a

hypothesis test for conditional independence with uniformly valid level $\alpha \in (0, 1)$ over \mathcal{H} ,

$$\sup_{P \in \mathcal{H}} E_P(\psi_n) \leq \alpha,$$

then the test cannot have power greater than α against any alternative $P \in \mathcal{Q} := \mathcal{P} \setminus \mathcal{H}$. This is true even when restricting the distribution of (X, Y, Z) to have bounded support. The purpose of this paper is to identify a subset $\mathcal{P}_0 \subset \mathcal{P}$ of distributions and a test ψ_n that has asymptotic (uniform) level over $\mathcal{P}_0 \cap \mathcal{H}$ and power against a large set of alternatives within $\mathcal{P}_0 \setminus \mathcal{H}$.

Our starting point is the so-called partial copula construction. Letting $F_{X|Z}$ and $F_{Y|Z}$ denote the conditional distribution functions of X given Z and Y given Z , respectively, we define random variables U_1 and U_2 by

$$U_1 := F_{X|Z}(X | Z) \quad \text{and} \quad U_2 := F_{Y|Z}(Y | Z).$$

Then the joint distribution of U_1 and U_2 is called the partial copula and it can be shown that $X \perp\!\!\!\perp Y | Z$ implies $U_1 \perp\!\!\!\perp U_2$. Thus the question about conditional independence can be transformed into a question about independence. The main challenge with this approach is that the conditional distribution functions are unknown and must be estimated.

In Section 3 we propose an estimator of conditional distribution functions based on quantile regression. More specifically, we let $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$ be a range of quantile levels for $0 < \tau_{\min} < \tau_{\max} < 1$, and let $Q(\mathcal{T} | z)$ denote the range of conditional \mathcal{T} -quantiles in the distribution $X | Z = z$. To estimate a conditional distribution function F given a sample $(X_i, Z_i)_{i=1}^n$ we propose to perform quantile regressions $\hat{q}_{k,z} = \hat{Q}^{(n)}(\tau_k | z)$ along an equidistant grid of quantile levels $(\tau_k)_{k=1}^m$ in \mathcal{T} , and then construct the estimator $\hat{F}^{(m,n)}$ by linear interpolation of the points $(\hat{q}_{k,z}, \tau_k)_{k=1}^m$. The main result of the first part of the paper is Theorem 5, which states that we can achieve the following bound on the estimation error

$$\|F - \hat{F}^{(m,n)}\|_{\mathcal{T}, \infty} := \sup_z \sup_{t \in Q(\mathcal{T}|z)} |F(t | z) - \hat{F}^{(m,n)}(t | z)| \in \mathcal{O}_P(g_P(n))$$

where g_P is a rate function describing the \mathcal{O}_P -consistency of the quantile regression procedure over the conditional \mathcal{T} -quantiles for P in a specified set of distributions $\mathcal{P}_0 \subset \mathcal{P}$. This result demonstrates how pointwise consistency of a quantile regression procedure over \mathcal{P}_0 can be transferred to the estimator $\hat{F}^{(m,n)}$, and we discuss how this can be extended to uniform consistency over \mathcal{P}_0 . We conclude the section by reviewing a flexible model class from quantile regression where such consistency results are available.

In Section 4 we describe a generic method for testing conditional independence based on estimated conditional distribution functions, $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$, obtained from a sample $(X_i, Y_i, Z_i)_{i=1}^n$. From these estimates we compute

$$\hat{U}_{1,i}^{(n)} := \hat{F}_{X|Z}^{(n)}(X_i | Z_i) \quad \text{and} \quad \hat{U}_{2,i}^{(n)} := \hat{F}_{Y|Z}^{(n)}(Y_i | Z_i),$$

for $i = 1, \dots, n$, which can then be plugged into a bivariate independence test. If $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ are consistent with a sufficiently fast rate of convergence, properties of the bivariate test, in terms of level and power, can be transferred to the test of conditional independence.

The details of this transfer of properties depend on the specific test statistic. The main contribution of the second part of the paper is a detailed treatment of a test given in terms of a generalized correlation, estimated as

$$\hat{\rho}_n := \frac{1}{n} \sum_{i=1}^n \varphi \left(\hat{U}_{1,i}^{(n)} \right) \varphi \left(\hat{U}_{2,i}^{(n)} \right)^T$$

for a function $\varphi = (\varphi_1, \dots, \varphi_q) : [0, 1] \rightarrow \mathbb{R}^q$ satisfying certain regularity conditions. A main result is Theorem 14, which states that $\sqrt{n}\hat{\rho}_n$ converges in distribution toward $\mathcal{N}(0, \Sigma \otimes \Sigma)$ under the hypothesis of conditional independence whenever $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ are \mathcal{O}_P -consistent with rates g_P and h_P satisfying $\sqrt{n}g_P(n)h_P(n) \rightarrow 0$. The covariance matrix Σ depends only on φ . We use this to show asymptotic pointwise level of the test when restricting to the set of distributions \mathcal{P}_0 where the required consistency can be obtained. We then proceed to show in Theorem 18 that $\sqrt{n}\hat{\rho}_n$ diverges in probability under a set of alternatives of conditional dependence when we have \mathcal{O}_P -consistency of the conditional distribution function estimators. This we use to show asymptotic pointwise power of the test. We also show how asymptotic uniform level and power can be achieved when $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ are uniformly consistent over \mathcal{P}_0 . Lastly, we provide an out-of-the-box procedure for conditional independence testing in conjunction with our quantile regression based conditional distribution function estimator $\hat{F}^{(m,n)}$ from Section 3.

In Section 5 we examine the proposed test through a simulation study where we assess the level and power properties of the test and benchmark it against existing nonparametric conditional independence tests. All proofs are collected in Appendix A.

2. Related Work

The partial copula and its application for conditional independence testing was initially introduced by Bergsma (2004) and further explored by Bergsma (2011). Its use for conditional independence testing has also been explored by Song (2009), Patra et al. (2016) and Liu et al. (2018). Moreover, properties of the partial copula was studied by Gijbels et al. (2015) and Spanhel and Kurz (2016) among others. A related but different approach for testing conditional independence via the factorization of the joint copula of (X, Y, Z) is given by Bouezmarni et al. (2012). Common for the existing approaches to using the partial copula for conditional independence testing is that the conditional distribution functions $F_{X|Z}$ and $F_{Y|Z}$ are estimated using a kernel smoothing procedure (Stute et al., 1986; Einmahl and Mason, 2005). The advantage of the approach is that the estimator is nonparametric, however, it does not scale well with the dimension of the conditioning variable Z . This is partly remedied by Haff and Segers (2015) who suggest a nonparametric pair-copula estimator whose convergence rate is independent of the dimension of Z . This estimator requires the simplifying assumption, which is a strong assumption not required for the validity of our approach. Moreover, it is not obvious how to incorporate parametric assumptions, such as a certain functional dependence between response and covariates, using kernel smoothing estimators, since there is only the choice of a kernel and a bandwidth. Furthermore, a treatment of the relationship between level and power properties of a partial copula based conditional independence test, and consistency of the conditional distribution function esti-

mator is lacking in the existing literature. In this work we take a novel approach to testing conditional independence using the partial copula by using quantile regression for estimating the conditional distribution functions. This allows for a distribution free modeling of the conditional distributions $X | Z = z$ and $Y | Z = z$ that can handle high-dimensionality of Z through penalization, and complicated response-predictor relationships by basis expansions. We also make the requirements on consistency of the conditional distribution function estimator that are needed to obtain level and power of the test explicit. A similar recent approach to testing conditional independence using regression methods is given by Shah and Peters (2020), who propose to test for vanishing correlation between the residuals after nonparametric conditional mean regression of X on Z and Y on Z . See also Ramsey (2014) and Fan et al. (2020). This approach captures dependence between X and Y given Z that lies in the conditional correlation. However, as is demonstrated through a simulation study in Section 5.5, it does not adequately account for conditional variance heterogeneity between X and Y given Z , while our partial copula based test captures the dependence more efficiently.

3. Estimation of Conditional Distribution Functions

Throughout the paper we restrict ourselves to the set of distributions \mathcal{P} over the hypercube $[0, 1]^{2+d}$ that are absolutely continuous with respect to Lebesgue measure. Let $(X, Y, Z) \sim P \in \mathcal{P}$ such that $X, Y \in [0, 1]$ and $Z \in [0, 1]^d$. When we speak of the distribution of X given Z relative to P we mean the conditional distribution that is induced when $(X, Y, Z) \sim P$. In this section we consider estimation of the conditional distribution function $F_{X|Z}$ of X given Z using quantile regression. Estimation of $F_{Y|Z}$ can be carried out analogously.

3.1 Conditional distribution and quantile functions

Given $z \in [0, 1]^d$ we denote by

$$F_{X|Z}(t | z) := P(X \leq t | Z = z)$$

the conditional distribution function of $X | Z = z$ for $t \in [0, 1]$. We denote by

$$Q_{X|Z}(\tau | z) := \inf\{t \in [0, 1] | F_{X|Z}(t | z) \geq \tau\}$$

the conditional quantile function of the conditional distribution $X | Z = z$ for $\tau \in [0, 1]$ and $z \in [0, 1]^d$. We will omit the subscript in $F_{X|Z}$ and $Q_{X|Z}$ when the conditional distribution of interest is clear from the context.

In quantile regression one models the function $z \mapsto Q(\tau | z)$ for fixed $\tau \in [0, 1]$. Estimation of the quantile regression function is carried out by solving the empirical risk minimization problem

$$\hat{Q}(\tau | \cdot) \in \arg \min_{f \in \mathcal{F}} \sum_{i=1}^n L_\tau(X_i - f(Z_i))$$

where the loss function $L_\tau(u) = u(\tau - 1(u < 0))$ is the so-called check function and \mathcal{F} is some function class. For $\tau = 1/2$ the loss function is $L_{1/2}(u) = |u|$, and we recover median regression as a special case. One can also choose to add regularization as with conditional mean regression. See Koenker (2005) and Koenker et al. (2017) for an overview of the field.

3.2 Quantile regression based estimator

Based on the conditional quantile function Q we define an approximation $\tilde{F}^{(m)}$ of the conditional distribution function F as follows. We let τ_{\min} and τ_{\max} denote fixed quantile levels satisfying $0 < \tau_{\min} < \tau_{\max} < 1$, and we let $q_{\min,z} := Q(\tau_{\min} | z) > 0$ and $q_{\max,z} := Q(\tau_{\max} | z) < 1$ denote the corresponding conditional quantiles.

Let $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$ denote the set of potential quantile levels. A grid in \mathcal{T} is a sequence $(\tau_k)_{k=1}^m$ such that $\tau_{\min} = \tau_1 < \dots < \tau_m = \tau_{\max}$ for $m \geq 2$. An equidistant grid is a grid $(\tau_k)_{k=1}^m$ for which $\tau_{k+1} - \tau_k$ is constant for $k = 1, \dots, m-1$. Also let $\tau_0 = 0$ and $\tau_{m+1} = 1$ be fixed.

Given a grid $(\tau_k)_{k=1}^m$ we let $q_{k,z} := Q(\tau_k | z)$ for $k = 1, \dots, m$ and define $q_{0,z} := 0$ and $q_{m+1,z} := 1$. For each $z \in [0, 1]^d$ we define a function $\tilde{F}^{(m)}(\cdot | z) : [0, 1] \rightarrow [0, 1]$ by linear interpolation of the points $(q_{k,z}, \tau_k)_{k=0}^{m+1}$:

$$\tilde{F}^{(m)}(t | z) := \sum_{k=0}^m \left(\tau_k + (\tau_{k+1} - \tau_k) \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} \right) 1_{(q_{k,z}, q_{k+1,z}]}(t). \quad (1)$$

Let $Q(\mathcal{T} | z) = [q_{\min,z}, q_{\max,z}]$ be the range of conditional \mathcal{T} -quantiles in the conditional distribution $X | Z = z$ for $z \in [0, 1]^d$, and define the supremum norm

$$\|f\|_{\mathcal{T}, \infty} = \sup_{z \in [0, 1]^d} \sup_{t \in Q(\mathcal{T} | z)} |f(t, z)|$$

for a function $f : [0, 1] \times [0, 1]^d \rightarrow \mathbb{R}$. Note that this is a norm on the set of bounded functions on $\{(t, z) | z \in [0, 1]^d, t \in Q(\mathcal{T} | z)\}$. Then we have the following approximation result.

Proposition 1 *Denote by $\tilde{F}^{(m)}$ the function (1) defined from a grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Then it holds that*

$$\|F - \tilde{F}^{(m)}\|_{\mathcal{T}, \infty} \leq \kappa_m$$

where $\kappa_m := \max_{k=1, \dots, m-1} (\tau_{k+1} - \tau_k)$ is the coarseness of the grid.

Choosing a finer and finer grid yields $\kappa_m \rightarrow 0$, which implies that $\tilde{F}^{(m)} \rightarrow F$ in the norm $\|\cdot\|_{\mathcal{T}, \infty}$ for $m \rightarrow \infty$.

By an estimator of the conditional distribution function F we mean a mapping from a sample $(X_i, Z_i)_{i=1}^n$ to a function $\hat{F}^{(n)}(\cdot | z) : [0, 1] \rightarrow [0, 1]$ such that for every $z \in [0, 1]^d$ it holds that $t \mapsto \hat{F}^{(n)}(t | z)$ is continuous and increasing with

$$\hat{F}^{(n)}(0 | z) = 0 \quad \text{and} \quad \hat{F}^{(n)}(1 | z) = 1.$$

Motivated by (1) we define the following estimator of the conditional distribution function.

Definition 2 *Let $(\tau_k)_{k=1}^m$ be a grid in \mathcal{T} . Define $\hat{q}_{0,z}^{(n)} := 0$ and $\hat{q}_{m+1,z}^{(n)} := 1$, and let $\hat{q}_{k,z}^{(n)} := \hat{Q}^{(n)}(\tau_k | z)$ for $k = 1, \dots, m$ be the predictions of a quantile regression model obtained from an i.i.d. sample $(X_i, Z_i)_{i=1}^n$. We define the estimator $\hat{F}^{(m,n)}(\cdot | z) : [0, 1] \rightarrow [0, 1]$ by*

$$\hat{F}^{(m,n)}(t | z) := \sum_{k=0}^m \left(\tau_k + (\tau_{k+1} - \tau_k) \frac{t - \hat{q}_{k,z}^{(n)}}{\hat{q}_{k+1,z}^{(n)} - \hat{q}_{k,z}^{(n)}} \right) 1_{(\hat{q}_{k,z}^{(n)}, \hat{q}_{k+1,z}^{(n)}]}(t) \quad (2)$$

for each $z \in [0, 1]^d$.

Note that the estimator is not monotone in the presence of quantile crossing (He, 1997). In this case we perform a re-arrangement of the estimated conditional quantiles in order to obtain monotonicity for finite sample size (Chernozhukov et al., 2010). However, the estimated conditional quantiles will be ordered correctly under the consistency assumptions that we will introduce in Assumption 1, that is, the re-arrangement becomes unnecessary, and the estimator becomes monotone with high probability as $n \rightarrow \infty$ for any grid $(\tau_k)_{k=1}^m$ in \mathcal{T} .

3.3 Pointwise consistency of $\hat{F}^{(m,n)}$

We will now demonstrate how pointwise consistency of the proposed estimator over a set of distributions $\mathcal{P}_0 \subset \mathcal{P}$ can be obtained under the assumption that the quantile regression procedure is pointwise consistent over \mathcal{P}_0 .

We will evaluate the consistency of $\hat{F}^{(m,n)}$ according to the supremum norm $\|\cdot\|_{\mathcal{T},\infty}$ introduced in Section 3.2, that is, we restrict the supremum to be over $t \in Q(\mathcal{T} | z)$ and not the entire interval $[0, 1]$. We do so because quantile regression generally does not give uniform consistency of all extreme quantiles, and in Section 4 we show how consistency of $\hat{F}^{(m,n)}$ between the conditional τ_{\min} - and τ_{\max} -quantiles is sufficient for conditional independence testing.

First, we have the following key corollary of Proposition 1, which is a simple application of the triangle inequality.

Corollary 3 *Let $\tilde{F}^{(m)}$ and $\hat{F}^{(m,n)}$ be given by (1) and (2), respectively. Then*

$$\|F - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty} \leq \kappa_m + \|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty}$$

for all grids $(\tau_k)_{k=1}^m$ in \mathcal{T} .

The random part of the right hand side of the inequality is the term $\|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty}$, while κ_m is deterministic and only depends on the choice of grid $(\tau_k)_{k=1}^m$. Controlling the term $\|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty}$ is an easier task than controlling $\|F - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty}$ directly because $\tilde{F}^{(m)}$ and $\hat{F}^{(m,n)}$ are piecewise linear, while F is only assumed to be continuous and increasing.

Consistency assumptions on the quantile regression procedure will allow us to show consistency of the estimator $\hat{F}^{(m,n)}$. Let the random variable

$$\mathcal{D}_{\mathcal{T}}^{(n)} := \sup_{z \in [0,1]^d} \sup_{\tau \in \mathcal{T}} |Q(\tau | z) - \hat{Q}^{(n)}(\tau | z)|$$

denote the uniform prediction error of a fitted quantile regression model, $\hat{Q}^{(n)}$, over the set of quantile levels $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$. Below we write $X_n \in \mathcal{O}_P(a_n)$ when X_n is big-O in probability of a_n with respect to P . See Appendix B for the formal definition.

Assumption 1 For each $P \in \mathcal{P}_0$ there exist

- (i) a deterministic rate function g_P tending to zero as $n \rightarrow \infty$ such that $\mathcal{D}_{\mathcal{T}}^{(n)} \in \mathcal{O}_P(g_P(n))$
- (ii) and a finite constant C_P such that the conditional density $f_{X|Z}$ satisfies

$$\sup_{x \in [0,1]} f_{X|Z}(x | z) \leq C_P$$

for almost all $z \in [0, 1]^d$.

Assumption 1 (i) is clearly necessary to achieve consistency of the estimator. Assumption 1 (ii) is a regularity condition that is used to ensure that $q_{k+1,z} - q_{k,z}$ does not tend to zero too fast as $\kappa_m \rightarrow 0$. We now have:

Proposition 4 Let Assumption 1 be satisfied. Then

$$\|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty} \in \mathcal{O}_P(g_P(n))$$

for each fixed $P \in \mathcal{P}_0$ and all equidistant grids $(\tau_k)_{k=1}^m$ in \mathcal{T} .

Consider letting the number of grid points m_n depend on the sample size n . By combining Corollary 3 and Proposition 4 we obtain the main pointwise consistency result.

Theorem 5 Let Assumption 1 be satisfied. Then

$$\|F - \hat{F}^{(m_n,n)}\|_{\mathcal{T},\infty} \in \mathcal{O}_P(g_P(n))$$

for each fixed $P \in \mathcal{P}_0$ given that the equidistant grids $(\tau_k)_{k=1}^{m_n}$ in \mathcal{T} satisfy $\kappa_{m_n} \in o(g_P(n))$.

This shows that $\hat{F}^{(m_n,n)}$ is pointwise consistent over \mathcal{P}_0 given that the quantile regression procedure is pointwise consistent over \mathcal{P}_0 . Moreover, we can transfer the rate of convergence g_P directly. In Section 4.4 we will use this type of pointwise consistency to show asymptotic pointwise level and power of our conditional independence test over \mathcal{P}_0 .

Note that we can estimate conditional distribution functions in settings with high dimensional covariates to the extent that the quantile regression estimation procedure can deal with high dimensionality. An example of such a procedure is given in Section 3.5.

We chose to state Proposition 4 and Theorem 5 for equidistant grids only, but in the proof of Proposition 4 we only need that the ratio κ_m/γ_m between the coarseness κ_m and the smallest subinterval $\gamma_m = \min_{k=1,\dots,m-1}(\tau_{k+1} - \tau_k)$ must not diverge as $m \rightarrow \infty$. This is obviously ensured for an equidistant grid. Moreover, for an equidistant grid, $\kappa_m = (\tau_{\max} - \tau_{\min})/(m - 1)$, and $\kappa_{m_n} \in o(g_P(n))$ if m_n grows with rate at least $g_P(n)^{-(1+\varepsilon)}$ for some $\varepsilon > 0$. Since the rate is unknown in practical applications we choose m to be the smallest integer larger than \sqrt{n} as a rule of thumb, since this represents the optimal parametric rate.

3.4 Uniform consistency of $\hat{F}^{(m,n)}$

The pointwise consistency result of Theorem 5 can be extended to a uniform consistency over \mathcal{P}_0 by strengthening Assumption 1 to hold uniformly. Below we write $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$ when X_n is big-O in probability of a_n uniformly over a set of distributions \mathcal{M} . We refer to Appendix B for the formal definition.

Assumption 2 *For $\mathcal{P}_0 \subset \mathcal{P}$ there exist*

- (i) *a deterministic rate function g tending to zero as $n \rightarrow \infty$ such that $\mathcal{D}_{\mathcal{T}}^{(n)} \in \mathcal{O}_{\mathcal{P}_0}(g(n))$*
- (ii) *and a finite constant C such that the conditional density $f_{X|Z}$ satisfies*

$$\sup_{x \in [0,1]} f_{X|Z}(x | z) \leq C$$

for almost all $z \in [0, 1]^d$.

With this stronger assumption we have a uniform extension of Proposition 4.

Proposition 6 *Let Assumption 2 be satisfied. Then*

$$\|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty} \in \mathcal{O}_{\mathcal{P}_0}(g(n)).$$

for all equidistant grids $(\tau_k)_{k=1}^m$ in \mathcal{T} .

We can now combine Corollary 3 with the stronger Proposition 6 to obtain the following uniform consistency of the estimator $\hat{F}^{(m,n)}$.

Theorem 7 *Suppose that Assumption 2 is satisfied. Then*

$$\|F - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty} \in \mathcal{O}_{\mathcal{P}_0}(g(n))$$

given that the equidistant grids $(\tau_k)_{k=1}^{m_n}$ in \mathcal{T} satisfy $\kappa_{m_n} \in o(g(n))$.

This shows that our estimator $\hat{F}^{(m,n)}$ can achieve uniform consistency over a set of distributions $\mathcal{P}_0 \subset \mathcal{P}$ given that the quantile regression procedure is uniformly consistent over \mathcal{P}_0 . In Section 4.5 we show how this strengthened result can be used to establish asymptotic uniform level and power of our conditional independence test over \mathcal{P}_0 .

3.5 A quantile regression model

In this section we will provide an example of a flexible quantile regression model and estimation procedure where consistency results are available. Consider the model

$$Q(\tau | z) = h(z)^T \beta_{\tau} \tag{3}$$

where $h : [0, 1]^d \rightarrow \mathbb{R}^p$ is a known and continuous transformation of Z , e.g., a polynomial or spline basis expansion to model non-linear effects. Inference in the model (3) was analyzed by Belloni and Chernozhukov (2011) and Belloni et al. (2019) in the high-dimensional

setup $p \gg n$. In the following we describe a subset of their results that is relevant for our application. Given an i.i.d. sample $(X_i, Z_i)_{i=1}^n$ and a fixed quantile regression level $\tau \in (0, 1)$, estimation of $\beta_\tau \in \mathbb{R}^p$ is carried out by penalized regression:

$$\hat{\beta}_\tau \in \arg \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^n L_\tau(X_i - h(Z_i)^T \beta) + \lambda_\tau \|\beta\|_1 \quad (4)$$

where $L_\tau(u) = u(\tau - 1(u < 0))$ is the check function, $\|\cdot\|_1$ is the 1-norm and $\lambda_\tau \geq 0$ is a tuning parameter that determines the degree of penalization. The tuning parameter λ_τ for a set \mathcal{Q} of quantile regression levels can be chosen in a data driven way as follows (Belloni and Chernozhukov, 2011, Section 2.3). Let $W_i = h(Z_i)$ denote the transformed predictors for $i = 1, \dots, n$. Then we set

$$\lambda_\tau = c\lambda\sqrt{\tau(1-\tau)} \quad (5)$$

where $c > 1$ is a constant with recommended value $c = 1.1$ and λ is the $(1 - n^{-1})$ -quantile of the random variable

$$\sup_{\tau \in \mathcal{T}} \frac{\|\Gamma^{-1} \frac{1}{n} \sum_{i=1}^n (\tau - 1(U_i \leq \tau) W_i)\|_\infty}{\sqrt{\tau(1-\tau)}}$$

where U_1, \dots, U_n are i.i.d. $\mathcal{U}[0, 1]$. Here $\Gamma \in \mathbb{R}^{p \times p}$ is a diagonal matrix with $\Gamma_{kk} = \frac{1}{n} \sum_{i=1}^n (W_i)_k^2$. The value of λ is determined by simulation.

Sufficient regularity conditions under which the above estimation procedure can be proven to be consistent are as follows.

Assumption 3 Denote by $f_{X|Z}$ the conditional density of X given Z . Let $c > 0$ and $C > 0$ be constants.

- (i) There exists s such that $\|\beta_\tau\|_0 \leq s$ for all $\tau \in \mathcal{Q} := [c, 1 - c]$.
- (ii) $f_{X|Z}$ is continuously differentiable such that $f_{X|Z}(Q_{X|Z}(\tau | z) | z) \geq c$ for each $\tau \in \mathcal{Q}$ and $z \in [0, 1]^d$. Moreover, $\sup_{x \in [0, 1]} f_{X|Z}(x | z) \leq C$ and $\sup_{x \in [0, 1]} \partial_x f_{X|Z}(x | z) \leq C$.
- (iii) The transformed predictor $W = h(Z)$ satisfies $c \leq E((W^T \theta)^2) \leq C$ for all $\theta \in \mathbb{R}^p$ with $\|\theta\|_2 = 1$. Moreover, $(E(\|W\|_\infty^{2q}))^{1/(2q)} \leq M_n$ for some $q > 2$ where M_n satisfies

$$M_n^2 \leq \frac{\delta_n n^{1/2-1/q}}{s \sqrt{\log(p \vee n)}}$$

and δ_n is some sequence tending to zero.

Assumption 3 (i) is a sparsity assumption, (ii) is a regularity condition on the conditional distribution, while (iii) is an assumption on the predictors. Examples of distributions for which Assumption 3 is satisfied are given in Belloni and Chernozhukov (2011) Section 2.5. This includes location models with Gaussian noise and location-scale models with bounded covariates, which in our setup with $Z \in [0, 1]^d$ means all location-scale models.

The following result (Belloni and Chernozhukov, 2011, Section 2.6) regarding the estimator $\hat{\beta}_\tau$ then holds.

Theorem 8 *Assume that the tuning parameters $\{\lambda_\tau \mid \tau \in \mathcal{Q}\}$ have been chosen according to (5). Then*

$$\sup_{\tau \in \mathcal{Q}} \|\beta_\tau - \hat{\beta}_\tau\|_2 \in \mathcal{O}_P \left(\sqrt{\frac{s \log(p \vee n)}{n}} \right)$$

under Assumption 3.

As a corollary of this consistency result we have the following.

Corollary 9 *Let $\hat{Q}(\tau \mid z) = h(z)^T \hat{\beta}_\tau$ be the predicted conditional quantile using the estimator $\hat{\beta}_\tau$. Then*

$$\sup_{z \in [0,1]^d} \sup_{\tau \in \mathcal{Q}} |Q(\tau \mid z) - \hat{Q}^{(n)}(\tau \mid z)| \in \mathcal{O}_P \left(\sqrt{\frac{s \log(p \vee n)}{n}} \right)$$

under Assumption 3.

This shows that Assumption 1 is satisfied under the model (3) whenever Assumption 3 is satisfied with $\mathcal{T} \subset \mathcal{Q}$ and $\sqrt{s \log(p \vee n)/n} \rightarrow 0$, which is the key underlying assumption of Theorem 5. Note also that Assumption 1 (ii) is contained in Assumption 3 (ii). Theorem 8 and Corollary 9 can be extended to hold uniformly over $\mathcal{P}_0 \subset \mathcal{P}$ by assuming that the conditions of Assumption 3 hold uniformly over \mathcal{P}_0 . This then gives the statement of Assumption 2 that is required for Theorem 7.

4. Testing Conditional Independence

In this section we describe the conditional independence testing framework in terms of the so-called partial copula. As above we let $(X, Y, Z) \sim P \in \mathcal{P}$ such that $X, Y \in [0, 1]$ and $Z \in [0, 1]^d$ where \mathcal{P} are the distributions that are absolutely continuous with respect to Lebesgue measure on $[0, 1]^{2+d}$. Also let f denote a generic density function. We then say that X is conditionally independent of Y given Z if

$$f(x, y \mid z) = f(x \mid z)f(y \mid z)$$

for almost all $x, y \in [0, 1]$ and $z \in [0, 1]^d$. See e.g. Dawid (1979). In this case we write that $X \perp\!\!\!\perp_P Y \mid Z$, where we usually omit the dependence on P when there is no ambiguity. By $\mathcal{H} \subset \mathcal{P}$ we denote the subset of distributions for which conditional independence is satisfied, and we let $\mathcal{Q} := \mathcal{P} \setminus \mathcal{H}$ be the alternative of conditional dependence.

4.1 The partial copula

We can regard the conditional distribution function as a mapping $(t, z) \mapsto F(t \mid z)$ for $t \in [0, 1]$ and $z \in [0, 1]^d$. Assuming that this mapping is measurable, we define a new pair of random variables U_1 and U_2 by the transformations

$$U_1 := F_{X|Z}(X \mid Z) \quad \text{and} \quad U_2 := F_{Y|Z}(Y \mid Z).$$

This transformation is usually called the probability integral transformation or Rosenblatt transformation due to Rosenblatt (1952), where the transformation was initially introduced and the following key result was shown.

Proposition 10 *It holds that $U_\ell \sim \mathcal{U}[0, 1]$ and $U_\ell \perp\!\!\!\perp Z$ for $\ell = 1, 2$.*

Hence the transformation can be understood as a normalization, where marginal dependencies of X on Z and Y on Z are filtered away. The joint distribution of U_1 and U_2 has been termed the partial copula of X and Y given Z in the copula literature (Bergsma, 2011; Spanhel and Kurz, 2016). Independence in the partial copula relates to conditional independence in the following way.

Proposition 11 *If $X \perp\!\!\!\perp Y \mid Z$ then $U_1 \perp\!\!\!\perp U_2$.*

Therefore the question about conditional independence can be transformed into a question about independence. Note, however, that $U_1 \perp\!\!\!\perp U_2$ does not in general imply that $X \perp\!\!\!\perp Y \mid Z$. See Property 7 in Spanhel and Kurz (2016) for a counterexample

The variables U_ℓ were termed nonparametric residuals by Patra et al. (2016) due to the independence property $U_\ell \perp\!\!\!\perp Z$ which is analogous to the property of conventional residuals in additive Gaussian noise models. Note that the entire conditional distribution function is required in order to compute the nonparametric residual, while conventional residuals in additive noise models are computed using only the conditional expectation. In return, Proposition 10 provides the distribution of the nonparametric residuals without assuming any functional or distributional relationship between X (Y resp.) and Z , whereas the distribution of conventional residuals is not known without further assumptions. Moreover, the nonparametric residuals U_1 and U_2 are independent under conditional independence, while conventional residuals are only uncorrelated unless we make a Gaussian assumption, say.

4.2 Generic testing procedure

Suppose $(X_i, Y_i, Z_i)_{i=1}^n$ is a sample from $P \in \mathcal{P}_0$ where \mathcal{P}_0 is some subset of \mathcal{P} . Also let $\mathcal{H}_0 := \mathcal{P}_0 \cap \mathcal{H}$ and $\mathcal{Q}_0 := \mathcal{P}_0 \cap \mathcal{Q}$ be the distributions in \mathcal{P}_0 satisfying conditional independence and conditional dependence, respectively. Denote by

$$U_{1,i} := F_{X|Z}(X_i \mid Z_i) \quad \text{and} \quad U_{2,i} := F_{Y|Z}(Y_i \mid Z_i)$$

the nonparametric residuals for $i = 1, \dots, n$. Let $\psi_n : [0, 1]^{2n} \rightarrow \{0, 1\}$ denote a test for independence in a bivariate continuous distribution. The observed value of the test is

$$\Psi_n := \psi_n((U_{1,i}, U_{2,i})_{i=1}^n)$$

with $\Psi_n = 0$ indicating acceptance and $\Psi_n = 1$ rejection of the hypothesis. By Proposition 11 we then reject the hypothesis of conditional independence, $X \perp\!\!\!\perp Y \mid Z$, if $\Psi_n = 1$. However, in order to implement the test in practice, we will need to replace the conditional distribution functions $F_{X|Z}$ and $F_{Y|Z}$ by estimates.

Given some generic estimators of the conditional distribution functions we can formulate a generic version of the partial copula conditional independence test as follows.

Definition 12 Let $(X_i, Y_i, Z_i)_{i=1}^n$ be an i.i.d. sample from $P \in \mathcal{P}_0$. Also let ψ_n be a test for independence in a bivariate continuous distribution.

(i) Form the estimates $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ based on $(X_i, Y_i, Z_i)_{i=1}^n$.

(ii) Compute the estimated nonparametric residuals

$$\hat{U}_{1,i}^{(n)} := \hat{F}_{X|Z}^{(n)}(X_i | Z_i) \quad \text{and} \quad \hat{U}_{2,i}^{(n)} := \hat{F}_{Y|Z}^{(n)}(Y_i | Z_i)$$

for $i = 1, \dots, n$.

(iii) Let $\hat{\Psi}_n := \psi_n \left((\hat{U}_{1,i}^{(n)}, \hat{U}_{2,i}^{(n)})_{i=1}^n \right)$ and reject the hypothesis $X \perp\!\!\!\perp Y | Z$ of conditional independence if $\hat{\Psi}_n = 1$.

This generic version of the conditional independence test is analogous to the approach of Bergsma (2011), but here we emphasize the modularity of the testing procedure. Firstly, one can use any method for estimating conditional distribution functions. Secondly, any test for independence in a bivariate continuous distribution can be utilized.

We note that under the assumptions of Theorem 5 it holds that

$$|(\hat{U}_{1,i}^{(n)}, \hat{U}_{2,i}^{(n)}) - (U_{1,i}, U_{2,i})|_{\mathcal{T},1} \xrightarrow{P} 0$$

where $|u - v|_{\mathcal{T},1} = |u_1 - v_1|1(u_1, v_1 \in \mathcal{T}) + |u_2 - v_2|1(u_2, v_2 \in \mathcal{T})$. That is, each estimated pair of nonparametric residuals has the partial copula as asymptotic distribution – except perhaps on the fringe part of the unit square outside of $\mathcal{T} \times \mathcal{T}$. This is a priori only a marginal result for each i , but it suggests that tests based on the estimated residuals behave as if they were i.i.d. observations from the partial copula.

Once we have chosen the test for independence, ψ_n , we can establish rigorous results on the properties of the test over the space of hypotheses \mathcal{H}_0 and alternatives \mathcal{Q}_0 , but how exactly to transfer the consistency of the estimated residuals to results on level and power depends on the specific test statistic. We will in the following sections demonstrate this transfer for one particular class of test statistics.

4.3 Generalized measure of correlation

We will now introduce a generalized measure of correlation that will form the basis for an independence test between the nonparametric residuals U_1 and U_2 .

Definition 13 The generalized correlation, ρ , between U_1 and U_2 is defined in term of a multivariate function $\varphi = (\varphi_1, \dots, \varphi_q) : [0, 1] \rightarrow \mathbb{R}^q$ as

$$\rho = E_P(\varphi(U_1)\varphi(U_2)^T) \tag{6}$$

such that ρ is a $q \times q$ matrix with entries $\rho_{k\ell} = E_P(\varphi_k(U_1)\varphi_\ell(U_2))$ for $k, \ell = 1, \dots, q$.

We will assume that the function $\varphi = (\varphi_1, \dots, \varphi_q)$ defining the generalized correlation satisfies the following assumptions for the remainder of the paper.

Assumption 4

- (i) The support \mathcal{T}_k of each coordinate function φ_k is a compact subset of $(0, 1)$.
- (ii) Each coordinate function $\varphi_k : [0, 1] \rightarrow \mathbb{R}$ is Lipschitz continuous.
- (iii) $\int_0^1 \varphi_k(u) du = 0$ and $\int_0^1 \varphi_k(u)^2 du = 1$ for each $k = 1, \dots, q$.
- (iv) The coordinate functions $\varphi_1, \dots, \varphi_q$ are linearly independent.

Let us provide some intuition about the interpretation of the generalized correlation ρ and explain the role of the assumptions on φ in Assumption 4.

Each entry $\rho_{k\ell}$ can be interpreted as an expected conditional correlation, and it can be understood in terms of the partial and conditional copula (Patton, 2006). Let $C(u_1, u_2 | z) = F(U_1 \leq u_1, U_2 \leq u_2 | Z = z)$ denote the conditional copula of X and Y given $Z = z$. Then the partial copula is the expected conditional copula, i.e., $C_p(u_1, u_2) = E_P(C(u_1, u_2 | Z))$ (Spanhel and Kurz, 2016). The conditional generalized correlation, $\rho_{k\ell}(z)$, between X and Y given $Z = z$ can be expressed in terms of the conditional copula by

$$\rho_{k\ell}(z) := E_P(\varphi_k(U_1)\varphi_\ell(U_2) | Z = z) = \int \varphi_k(u_1)\varphi_\ell(u_2)C(du_1, du_2 | z).$$

By the tower property of conditional expectations, $\rho_{k\ell}$ can be represented as an expected generalized correlation

$$\rho_{k\ell} = E_P(\rho_{k\ell}(Z)) = \int \varphi_k(u_1)\varphi_\ell(u_2)C_p(du_1, du_2).$$

Hence $\rho_{k\ell}$ measures the expected conditional generalized correlation of X and Y given Z w.r.t. the distribution of Z . Importantly, Assumption 4 (iii) implies that

$$\rho = E_P(\varphi(U_1)\varphi(U_2)^T) = E_P(\varphi(U_1))E(\varphi(U_2)^T) = 0$$

whenever $X \perp\!\!\!\perp Y | Z$ due to Proposition 11.

The purpose of Assumption 4 (i) is twofold. Firstly, letting the supports \mathcal{T}_k and \mathcal{T}_ℓ of φ_k and φ_ℓ be subsets of $(0, 1)$ implies that $\rho_{k\ell}$ focuses on dependence in the compact region $\mathcal{T}_k \times \mathcal{T}_\ell \subset (0, 1)^2$ of the outcome space $[0, 1]^2$ of (U_1, U_2) . For $q \geq 2$ the generalized correlation ρ thus summarizes dependencies in different regions of the outcome space. See Figure 1 for an illustration of this idea. Secondly, the supports $(\mathcal{T}_k)_{k=1}^q$ will play the role as subsets of the possible quantile levels $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$, when we choose the conditional distribution function estimators based on quantile regression from Section 3.2. This connection will be made clear in Section 4.4.

The functional form of φ_k and φ_ℓ determines the kind of dependence measured by $\rho_{k\ell}$. Ignoring Assumption 4 (i), consider letting $\varphi_k(u) = \varphi_\ell(u) = \sqrt{12}(u - 1/2)$ for $u \in [0, 1]$. Then $\rho_{k\ell}$ measures the expected conditional Spearman correlation between X and Y given Z with respect to the distribution of Z . In Section 4.6 we describe a choice of functions φ_k that leads to a trimmed version of expected conditional Spearman correlation which satisfies Assumption 4 (i). As we shall see in Section 4.4, Assumption 4 (ii), i.e., that the coordinate functions φ_k are Lipschitz continuous, is crucial for deriving asymptotic properties for the empirical version of the generalized correlation ρ . Lastly, we assume that $\varphi_1, \dots, \varphi_q$ are linearly independent in Assumption 4 (iv) to avoid degeneracy of its empirical version, which we introduce in Section 4.4.

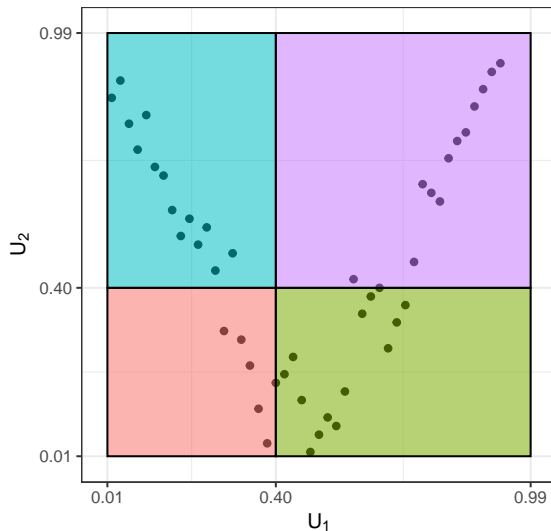


Figure 1: A sample from a copula (U_1, U_2) with clear dependence, but where the overall sample correlation is close to zero. The dependence is captured by considering sample correlation of observations in different regions of the outcome space.

4.4 Test based on generalized correlation

In this section we will analyze in depth the conditional independence test resulting from basing the test ψ_n in Definition 12 on the generalized correlation ρ . We will formulate the results in terms of a generic method for estimating conditional distribution functions in order to emphasize the generality of the method and illustrate the abstract assumptions needed for the test to be sound. Along the way we will explain when the assumptions are satisfied for the quantile regression based estimator $\hat{F}^{(m,n)}$ that we developed in Section 3.

With ρ the generalized correlation between U_1 and U_2 defined in terms of a function φ satisfying Assumption 4 we let $\rho_n : [0, 1]^{2n} \rightarrow \mathbb{R}^{q \times q}$ be its corresponding empirical version:

$$\rho_n((u_i, v_i)_{i=1}^n) := \frac{1}{n} \sum_{i=1}^n \varphi(u_i) \varphi(v_i)^T. \quad (7)$$

Soundness of a test based on ρ_n depends on consistency of the estimators $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$. Recall that we by $\mathcal{T}_1, \dots, \mathcal{T}_q$ denote the supports of $\varphi_1, \dots, \varphi_q$. Let $\tau_{\min} := \inf(\mathcal{T}_1 \cup \dots \cup \mathcal{T}_q) > 0$ and $\tau_{\max} := \sup(\mathcal{T}_1 \cup \dots \cup \mathcal{T}_q) < 1$, and then define $\mathcal{T} := [\tau_{\min}, \tau_{\max}]$. As in Section 3.2 we let the norm $\|\cdot\|_{\mathcal{T}, \infty}$ be given by

$$\|f(t, z)\|_{\mathcal{T}, \infty} = \sup_{z \in [0, 1]^d} \sup_{t \in Q_{X|Z}(\mathcal{T}|z)} |f(t, z)|$$

when X given Z is the conditional distribution of interest. Similarly define $\|\cdot\|'_{\mathcal{T}, \infty}$ by

$$\|f(t, z)\|'_{\mathcal{T}, \infty} = \sup_{z \in [0, 1]^d} \sup_{t \in Q_{Y|Z}(\mathcal{T}|z)} |f(t, z)|.$$

Then we have the following assumption on our estimators.

Assumption 5 *For each distribution $P \in \mathcal{P}_0$ there exist deterministic rate functions g_P and h_P tending to zero as $n \rightarrow \infty$ and functions $\xi, \xi' : [0, 1] \times [0, 1]^d \rightarrow \mathbb{R}$ such that*

$$(i) \quad \|F_{X|Z} - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}, \infty} \in \mathcal{O}_P(g_P(n)) \text{ and } \|F_{Y|Z} - \hat{F}_{Y|Z}^{(n)}\|'_{\mathcal{T}, \infty} \in \mathcal{O}_P(h_P(n)).$$

$$(ii) \quad \|\xi - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}^c, \infty} \in \mathcal{O}_P(g_P(n)) \text{ and } \|\xi' - \hat{F}_{Y|Z}^{(n)}\|'_{\mathcal{T}^c, \infty} \in \mathcal{O}_P(h_P(n)).$$

Assumption 5 (i) states that our estimators $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ are consistent with rates g_P and h_P over the conditional \mathcal{T} -quantiles in their respective conditional distributions. This is the result of Theorem 5 regarding our quantile regression based estimator $\hat{F}^{(m,n)}$ when \mathcal{T} as above is taken as the set of potential quantile regression levels.

Assumption 5 (ii) is an assumption on the behavior of our estimator in the tails of the conditional distribution, i.e., over the conditional \mathcal{T}^c -quantiles. Here we do not assume consistency, but we do assume that the limit for $n \rightarrow \infty$ exists, and that our estimators are convergent to their limits with rates g_P and h_P respectively. This assumption is satisfied by our quantile regression based estimator $\hat{F}^{(m,n)}$ whenever it satisfies Assumption 5 (i).

With this assumption we first establish the asymptotic distribution of the test statistic

$$\hat{\rho}_n := \rho_n \left((\hat{U}_{1,i}, \hat{U}_{2,i})_{i=1}^n \right) = \frac{1}{n} \sum_{i=1}^n \varphi(\hat{U}_{1,i}) \varphi(\hat{U}_{2,i})^T \quad (8)$$

under the hypothesis of conditional independence. Below we use \Rightarrow_P to denote convergence in distribution with respect to P .

Theorem 14 *Suppose that Assumption 5 is satisfied with rate functions g_P and h_P such that $\sqrt{n}g_P(n)h_P(n) \rightarrow 0$ as $n \rightarrow \infty$ for each $P \in \mathcal{P}_0$. Then the statistic $\hat{\rho}_n$ given by (8) satisfies*

$$\sqrt{n}\hat{\rho}_n \Rightarrow_P \mathcal{N}(0, \Sigma \otimes \Sigma)$$

for each fixed $P \in \mathcal{H}_0$. The asymptotic covariance matrix is given by

$$\Sigma_{k,s} = \int_0^1 \varphi_k(u) \varphi_s(u) du$$

for $k, s = 1, \dots, q$ and does not depend on P .

If the rate functions are $g_P(n) = n^{-a}$ and $h_P(n) = n^{-b}$, then we require that $a+b > 1/2$. Thus convergence slightly faster than rate $n^{-1/4}$ for both estimators is sufficient, but there can also be a tradeoff between the rates. Interestingly, Theorem 14 does not require sample splitting for the estimation of the conditional distribution function and computation of the test statistic. This is due to the fact that we are only interested in the asymptotic distribution under conditional independence. A similar phenomenon was found by Shah and Peters (2020), when they proved asymptotic normality of their Generalised Covariance Measure under conditional independence.

According to Corollary 9, the requirement $\sqrt{n}g_P(n)h_P(n) \rightarrow 0$ is satisfied for our quantile regression based estimator $\hat{F}^{(m,n)}$, if the quantile regression model (3) of Section 3.5 is valid for both X given Z and Y given Z for some continuous transformations $h_1 : [0, 1]^d \rightarrow \mathbb{R}^{p_1}$ and $h_2 : [0, 1]^d \rightarrow \mathbb{R}^{p_2}$ and Assumption 3 is satisfied with $s_1, s_2, p_1, p_2, n \rightarrow \infty$ such that

$$\sqrt{n} \cdot \sqrt{\frac{s_1 \log(p_1 \vee n)}{n}} \cdot \sqrt{\frac{s_2 \log(p_2 \vee n)}{n}} = \sqrt{\frac{s_1 s_2 \log(p_1 \vee n) \log(p_2 \vee n)}{n}} \rightarrow 0$$

where $s_1 = \sup_{\tau \in \mathcal{T}} \|\beta_{1,\tau}\|_0$ and $s_2 = \sup_{\tau \in \mathcal{T}} \|\beta_{2,\tau}\|_0$ are the sparsities of the model parameters.

With the test statistic

$$T_n := \|\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2}\|_F^2, \quad (9)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, we have the following corollary of Theorem 14.

Corollary 15 *Let the condition of Theorem 14 be satisfied and let T_n be given by (9). Then it holds that*

$$nT_n \Rightarrow_P \chi_{q^2}^2$$

for each fixed $P \in \mathcal{H}_0$.

In view of Theorem 14 and Corollary 15 we define the following conditional independence test based on the generalized correlation.

Definition 16 *Let $\alpha \in (0, 1)$ be a desired significance level and T_n the test statistic (9). Then we let $\hat{\Psi}_n$ be the test given by*

$$\hat{\Psi}_n = 1(T_n > n^{-1} z_{1-\alpha})$$

where $z_{1-\alpha}$ is the $(1 - \alpha)$ -quantile of a $\chi_{q^2}^2$ -distribution.

Control of the asymptotic pointwise level is then an easy corollary of Corollary 15.

Corollary 17 *Suppose that Assumption 5 is satisfied with rate functions g_P and h_P such that $\sqrt{n}g_P(n)h_P(n) \rightarrow 0$ as $n \rightarrow \infty$ for each $P \in \mathcal{P}_0$. Then the test $\hat{\Psi}_n$ given by Definition 16 has asymptotic pointwise level over \mathcal{H}_0 , i.e.,*

$$\limsup_{n \rightarrow \infty} E_P(\hat{\Psi}_n) = \alpha$$

for each fixed $P \in \mathcal{H}_0$.

This shows that the test achieves correct level given consistency of the estimators $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ with suitably fast rates. To obtain results on power of the test $\hat{\Psi}_n$ we only need to understand how $\hat{\rho}_n$ converges in probability and not its entire asymptotic distribution.

Theorem 18 *The test statistic $\hat{\rho}_n$ given by (8) satisfies*

$$\hat{\rho}_n \xrightarrow{P} \rho$$

for each fixed $P \in \mathcal{P}_0$ under Assumption 5.

One may note that the theorem does not require that the rate functions g_P and h_P converge to zero at a certain rate. Let $\mathcal{A}_0 \subseteq \mathcal{Q}_0$ be the subset of alternatives for which $\rho_{k\ell} \neq 0$ for at least one combination of $k, \ell = 1, \dots, q$. Then we have the following corollary of Theorem 18, which exploits that nT_n diverges in probability whenever $P \in \mathcal{A}_0$.

Corollary 19 *For each level $\alpha \in (0, 1)$ the test $\hat{\Psi}_n$ given by Definition 16 has asymptotic pointwise power against \mathcal{A}_0 , i.e.,*

$$\liminf_{n \rightarrow \infty} E_P(\hat{\Psi}_n) = 1$$

for each fixed $P \in \mathcal{A}_0$ under Assumption 5.

Let us discuss the alternatives the test has power against. Firstly, note that we always have the implications

$$X \perp\!\!\!\perp Y \mid Z \quad \Rightarrow \quad U_1 \perp\!\!\!\perp U_2 \quad \Rightarrow \quad \rho = 0$$

However, none of the reverse implications are in general true. We do, however, have the following result stating a sufficient condition for the reverse implication of the first statement.

Proposition 20 *Assume that $(U_1, U_2) \perp\!\!\!\perp Z$. Then $X \perp\!\!\!\perp Y \mid Z$ if and only if $U_1 \perp\!\!\!\perp U_2$.*

This means that if Z only affects the marginal distributions of X and Y , then independence in the partial copula implies conditional independence. This is known as the simplifying assumption in the copula literature (Gijbels et al., 2015; Spanhel and Kurz, 2015). Naturally, $U_1 \perp\!\!\!\perp U_2$ always implies $X \perp\!\!\!\perp Y \mid Z$, so the simplifying assumption is not a necessary condition for our test to have power, but it does give some intuition about a subset of distributions for which the partial copula completely characterizes conditional independence. However, an unavoidable limitation of the method is that it can never have power against alternatives for which $U_1 \perp\!\!\!\perp U_2$ but $X \not\perp\!\!\!\perp Y \mid Z$.

Turning to the second implication, Corollary 19 tells us that we have power against alternatives for which $\rho_{k\ell} \neq 0$ for some $k, \ell = 1, \dots, q$. However, not all types of dependencies can be detected in this fashion, and it is possible that $\rho = 0$, while $U_1 \not\perp\!\!\!\perp U_2$. A test based on ρ will not have power against such an alternative. For an abstract interpretation of the generalized correlation ρ we refer to Section 4.3. In Section 4.6 we introduce a concrete generalized correlation and elaborate on its interpretation.

Finally, basing the test on values of T_n is natural since the asymptotic behaviour is readily available through Theorem 14, but other transformations of $\hat{\rho}_n$ could be considered such as taking the coordinatewise absolute maximum $\max_{k,l} |(\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2})_{k,l}| = \|\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2}\|_\infty$.

4.5 Uniform level and power results

The level and power results of Section 4.4 are pointwise over the space of hypotheses and alternatives, i.e., they state level and power of the test when fixing a distribution P . In this section we describe how these results can be extended to hold uniformly by strengthening the statements in Assumption 5 to hold uniformly.

Assumption 6 For $\mathcal{P}_0 \subset \mathcal{P}$ there exist deterministic rate functions g and h tending to zero as $n \rightarrow \infty$ and functions $\xi, \xi' : [0, 1] \times [0, 1]^d \rightarrow \mathbb{R}$ such that

$$(i) \quad \|F_{X|Z} - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}, \infty} \in \mathcal{O}_{\mathcal{P}_0}(g(n)) \text{ and } \|F_{Y|Z} - \hat{F}_{Y|Z}^{(n)}\|'_{\mathcal{T}, \infty} \in \mathcal{O}_{\mathcal{P}_0}(h(n)).$$

$$(ii) \quad \|\xi - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}^c, \infty} \in \mathcal{O}_{\mathcal{P}_0}(g(n)) \text{ and } \|\xi' - \hat{F}_{Y|Z}^{(n)}\|'_{\mathcal{T}^c, \infty} \in \mathcal{O}_{\mathcal{P}_0}(h(n)).$$

As before we note that Assumption 6 (i) is the result of Theorem 7 regarding our quantile regression based estimator $\hat{F}^{(m,n)}$. Moreover, Assumption 6 (ii) is valid for $\hat{F}^{(m,n)}$ whenever it satisfies Assumption 6 (i).

We will now describe the extensions of Theorem 14 and Theorem 18 that can be obtained under Assumption 6. Below we write $\Rightarrow_{\mathcal{M}}$ to denote uniform convergence in distribution over a set of distributions \mathcal{M} , and we use $\rightarrow_{\mathcal{M}}$ to denote uniform convergence in probability over \mathcal{M} . We refer to Appendix B for the formal definitions.

Theorem 21 Let $\hat{\rho}_n$ be the statistic given by (8). Then we have:

(i) Under Assumption 6 with rate functions satisfying $\sqrt{n}g(n)h(n) \rightarrow 0$ it holds that

$$\sqrt{n}\hat{\rho}_n \Rightarrow_{\mathcal{H}_0} \mathcal{N}(0, \Sigma \otimes \Sigma)$$

where Σ is as in Theorem 14.

(ii) Under Assumption 6 it holds that $\hat{\rho}_n \rightarrow_{\mathcal{P}_0} \rho$.

As a straightforward corollary of Theorem 21 (i) we get the following uniform level result.

Corollary 22 The test $\hat{\Psi}_n$ given by Definition 16 has asymptotic uniform level over \mathcal{H}_0 , i.e.,

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{H}_0} E_P(\hat{\Psi}_n) = \alpha,$$

given that Assumption 6 is satisfied with $\sqrt{n}g(n)h(n) \rightarrow 0$ as $n \rightarrow \infty$.

The pointwise power result of Theorem 19 does not extend directly to a uniform version in the same way as the level result. For $\lambda > 0$ we let $\mathcal{A}_\lambda \subset \mathcal{Q}_0$ be the set of alternatives for which $|(\rho_P)_{k\ell}| > \lambda$ for at least one combination of $k, \ell = 1, \dots, q$, where we emphasize that ρ_P depends on the distribution P . We then have the following uniform power result as a corollary of Theorem 21 (ii).

Corollary 23 For all fixed levels $\alpha \in (0, 1)$ the test $\hat{\Psi}_n$ given by Definition 16 has asymptotic uniform power against \mathcal{A}_λ for each $\lambda > 0$, i.e.,

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{A}_\lambda} E_P(\hat{\Psi}_n) = 1,$$

under Assumption 6.

The reason we need to restrict to the class of alternatives \mathcal{A}_λ for a fixed $\lambda > 0$ is the following. If the infimum is taken over \mathcal{A}_0 , then there could exist a sequence $(P_m)_{m=1}^\infty \subset \mathcal{A}_0$ of distributions such that $(\rho_{P_m})_{k\ell} \neq 0$ for each $m \geq 1$ but $(\rho_{P_m})_{k\ell} \rightarrow 0$ as $m \rightarrow \infty$. As a consequence nT_n will not necessarily diverge in probability, which is crucial to the proof of the corollary. However, when restricting to \mathcal{A}_λ we are ensured that $\inf_{P \in \mathcal{A}_\lambda} |(\rho_P)_{k\ell}| \geq \lambda > 0$ for at least one combination of $k, \ell = 1, \dots, q$.

We note that these uniform level and power results are not in contradiction with the impossibility result of Shah and Peters (2020) mentioned in Section 1 because our results apply to a restricted set of distributions, \mathcal{P}_0 , where the conditional distribution functions are estimable with sufficiently fast rate.

4.6 Trimmed Spearman correlation

We will now define a specific family of functions φ defining the generalized correlation that can be shown to satisfy Assumption 4, which results in trimmed versions of the expected conditional Spearman correlation. As mentioned in Section 4.3, ignoring Assumption 4 (i), we could consider

$$\varphi_k(u) = \varphi_\ell(u) = \sqrt{12} \left(u - \frac{1}{2} \right) \tag{10}$$

for $u \in [0, 1]$ which results in $\rho_{k\ell}$ being the expected conditional Spearman correlation of X and Y given Z with respect to the distribution of Z . Drawing inspiration from (10) we define a class of functions $\varphi : [0, 1] \rightarrow \mathbb{R}^q$ by letting

$$\varphi_k(u) = c_k(u - m_k)\sigma_k(u) \tag{11}$$

such that each $\varphi_k : [0, 1] \rightarrow \mathbb{R}$ is determined by a Lipschitz continuous function $\sigma_k : [0, 1] \rightarrow \mathbb{R}$ with the support \mathcal{T}_k of σ_k a compact interval in $(0, 1)$, $\int_0^1 \sigma_k(u) du = 1$ and

$$m_k = \int u \sigma_k(u) du \quad \text{and} \quad c_k = \left(\int (u - m_k)^2 \sigma_k(u)^2 du \right)^{-1/2}.$$

The choice (11) satisfies Assumption 4 (i) – (iii) by construction, and if e.g. $\mathcal{T}_k \setminus \cup_{\ell \neq k} \mathcal{T}_\ell \neq \emptyset$ the functions are also linearly independent. We call the resulting generalized correlation ρ a trimmed Spearman correlation, and we refer to the functions σ_k as trimming functions. Note that if the supports $(\mathcal{T}_k)_{k=1}^q$ of $(\sigma_k)_{k=1}^q$ are chosen to be disjoint, then the covariance matrix Σ of Theorem 14 is the identity matrix.

A starting point for choosing a trimming function σ is the normalized indicator

$$u \mapsto (\lambda - \mu)^{-1} \mathbf{1}_{[\mu, \lambda]}(u) \tag{12}$$

for $u \in [0, 1]$ where $0 < \mu < \lambda < 1$ are trimming parameters. However, (12) is not a valid trimming function, since it is not Lipschitz. Therefore, we consider a simple linear approximation $\sigma : [0, 1] \rightarrow \mathbb{R}$ given by

$$\sigma(u) = Kf(u) \quad \text{and} \quad f(u) = \begin{cases} 1, & u \in [\mu + \delta, \lambda - \delta] \\ 0, & u \in [\mu, \lambda]^c \\ \delta^{-1}(u - \mu), & u \in [\mu, \mu + \delta) \\ \delta^{-1}(\lambda - u), & u \in (\lambda - \delta, \lambda] \end{cases} \quad (13)$$

and $K = (\lambda - \mu - \delta)^{-1}$. Here $0 < \delta < (\lambda - \mu)/2$ is a fixed parameter that determines the accuracy of the approximation. It is elementary to verify that σ given by (13) is a valid trimming function, i.e., σ is Lipschitz continuous with $\int \sigma(u)du = 1$ and support $[\mu, \lambda] \subset (0, 1)$.

The interpretation of a generalised correlation ρ based on φ of the form (11) with trimming function σ of the form (13) is as follows. The entry $\rho_{k\ell}$ is an approximation of the expected conditional Spearman correlation between the observations of X and Y , respectively, that lie in the \mathcal{T}_k -quantile range of the distribution of X given Z and the \mathcal{T}_ℓ -quantile range of the distribution of Y given Z , respectively, with respect to the distribution of Z . The matrix ρ then summarizes this type of dependence within different quantile ranges of X and Y given Z .

4.7 Practical considerations

Throughout Sections 4.4 and 4.5 we have analyzed our proposed test for conditional independence with an emphasis on modularity of the method regarding the choice of estimators $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ of the conditional distribution functions and the choice of the function φ that defines the generalized correlation ρ of Section 4.3. This focus on the conceptual assumptions displays the generality of the method, but it also leaves the practitioner of conditional independence testing with a number of choices to be made. In this section we summarize a set of choices to make the method work out-of-the-box.

Throughout the paper we have assumed that all random variables take values in the unit interval, i.e., $(X, Y, Z) \in [0, 1]^{d+2}$. This is not a restriction, since if e.g. $X \in \mathbb{R}$ we can always apply a strictly increasing, continuous transformation $t : \mathbb{R} \rightarrow [0, 1]$ to obtain a new random variable $X' = t(X)$ with values in $[0, 1]$. Moreover, the initial conditional independence structure of (X, Y, Z) is preserved since the transformation is marginal on X and bijective. The transformation t can be chosen to be e.g. the logistic function.

In principle, an arbitrary and fixed marginal transformation could be used for all variables, but we recommend to transform data to the unit interval via marginal empirical distribution functions. This results in transformed variables known in the copula literature as pseudo copula observations. The transformation creates dependence, similar to the dependence created by other common preprocessing techniques such as centering and scaling, which the theoretical analysis has not accounted for. We suggest, nevertheless, to use this preprocessing technique in practise, and in the simulation study in Section 5 we use pseudo copula observations since it reflects how a practitioner would transform the variables.

To estimate the conditional distribution functions $\hat{F}_{X|Z}^{(m,n)}$ and $\hat{F}_{Y|Z}^{(m,n)}$ using Definition 2, we suggest choosing $\tau_{\min} = 0.01$ and $\tau_{\max} = 0.99$ and form the equidistant grid $(\tau_k)_{k=1}^m$ in $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$ with the number of gridpoints $m = \lceil \sqrt{n} \rceil$. We then suggest using a model of the form (3) for both the quantile regression model $Q_{X|Z}(\tau_k | \cdot)$ and $Q_{Y|Z}(\tau_k | \cdot)$ for each $k = 1, \dots, m$, where the bases h_1 and h_2 can be chosen to be e.g. an additive B-spline basis for each component of Z .

To test the hypothesis of conditional independence we suggest using the $\hat{\Psi}_n$ from Definition 16 based on the estimated nonparametric residuals $(\hat{U}_{1,i}, \hat{U}_{2,i})_{i=1}^n$. To this end we choose $q \geq 1$ and let $\tau_{\min} = \lambda_0 < \dots < \lambda_q = \tau_{\max}$ be an equidistant grid in \mathcal{T} . We then define the trimming function σ_k to have the form (13) with trimming parameters λ_k and λ_{k+1} and approximation parameter $\delta = 0.01 \cdot (\lambda_{k+1} - \lambda_k)$ for each $k = 0, \dots, q - 1$. We then define $(\sigma_k)_{k=1}^q$ according to (11), compute the test statistic $\hat{\rho}_n$ using (8) and compute $\hat{\Psi}_n$ as in Definition 16 using a desired significance level $\alpha \in (0, 1)$.

There are two non-trivial choices remaining. The first is the choice of bases h_1 and h_2 for the quantile regression models $Q_{X|Z}(\tau_k | z) = h_1(z)^T \beta_{\tau_k}$ and $Q_{Y|Z}(\tau_k | z) = h_2(z)^T \beta_{\tau_k}$. Here the practitioner needs to make a qualified model selection. We recommend using a flexible basis such as an additive B-spline basis, and perform penalized estimation using (4) to avoid overfitting. The second choice is the dimension of the generalized correlation $q \geq 1$, which corresponds to a choice of independence test in the partial copula. Note that the generalized correlation ρ as above is defined for any $q \geq 1$, and there is conditional dependence, $X \not\perp\!\!\!\perp Y | Z$, if there exists $q \geq 1$ for which $\rho \neq 0$. We suggest trying one or a few, small values, e.g. $q \in \{1, \dots, 5\}$, and reject the hypothesis of conditional independence if one of the tests rejects the hypothesis, but of course be aware of multiple testing issues.

5. Simulation Study

In this section we examine the performance of our conditional independence test $\hat{\Psi}_n$ of Definition 16, when combining it with the quantile regression based conditional distribution function estimator $\hat{F}^{(m,n)}$ from Definition 2. Firstly, we verify the level and power results obtained in Section 4.4 and Section 4.5 empirically. Secondly, we compare the test with other conditional independence tests. The test was implemented in the R language (R Core Team, 2021) using the `quantreg` package (Koenker, 2021) as the backend for performing quantile regression. The implementation and code for producing the simulations can be obtained from <https://github.com/lassepetersen/partial-copula-CI-test>.

5.1 Evaluation method

We will evaluate the tests by their ability to hold level when data is generated from a distribution where conditional independence holds, and by their power when data is generated from a distribution where conditional independence does not hold. In order to make the results independent of a chosen significance level we will base the evaluation on the p -values of the tests.

If a test has valid level, then we expect the p -value to be asymptotically $\mathcal{U}[0, 1]$ -distributed. In Sections 5.3 and 5.4 we evaluate the level by a Kolmogorov-Smirnov (KS) statistics as a function of sample size n , which is independent of any specific significance level. A small KS

statistic is an indication of valid level. To examine power we consider in Sections 5.3 and 5.4 the p -values of the test as a function of the sample size, where we expect the p -values to tend to zero under the alternative of conditional dependence. Here a small p -value is an indication of large power. In Section 5.5 we evaluate the power against a local alternative, which shrinks with the sample size n toward the hypothesis of conditional independence with rate $n^{-\frac{1}{2}}$.

5.2 Data generating processes

This section gives an overview of the data generating processes that we use for benchmarking and comparison. The first category consists of data generating processes of the form

$$X = f_1(Z) + g_1(Z) \cdot \varepsilon_1 \quad \text{and} \quad Y = f_2(Z) + g_2(Z) \cdot \varepsilon_2 \quad (\text{H})$$

where $f_1, f_2, g_1, g_2 : \mathbb{R}^d \rightarrow \mathbb{R}$ belong to some function class and $\varepsilon_1, \varepsilon_2$ are independent errors. For data generating processes of type (H), conditional independence is satisfied. The second category consists of data generating processes of the form

$$X = f_1(Z) + g_1(Z) \cdot \varepsilon_1 \quad \text{and} \quad Y = f_2(Z, X) + g_2(Z, X) \cdot \varepsilon_2 \quad (\text{A})$$

where again $f_1, g_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ and $f_2, g_2 : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ belong to some function class and $\varepsilon_1, \varepsilon_2$ are independent errors. Under data generating processes of type (A), conditional independence is not satisfied. We will consider four different data generating processes corresponding to different choices of functions f_1, g_1, f_2 and g_2 and error distributions.

- (1) For data generating processes H_1 and A_1 we let

$$f_k(w_1, \dots, w_d) = \sum_{j=1}^d \beta_{1,k,j} w_j + \beta_{2,k,j} w_j^2$$

$$g_k(w_1, \dots, w_d) = \exp \left(- \left| \sum_{j=1}^d \alpha_{1,k,j} w_j + \alpha_{2,k,j} w_j^2 \right| \right)$$

for $k = 1, 2$ and real valued coefficients $(\alpha_{\ell,k,j}, \beta_{\ell,k,j})_{\ell=1,2, k=1,2, j=1, \dots, d}$. Here each $Z_j \sim \mathcal{U}[-1, 1]$ independently, ε_1 follows an asymmetric Laplace distribution with location 0, scale 1 and skewness 0.8, and ε_2 follows a Gumpel distribution with location 0 and scale 1.

- (2) For data generating processes H_2 and A_2 we let $g_1 = g_2 = 1$ and

$$f_k(w_1, \dots, w_d) = \sum_{j=1}^d \beta_{k,j} w_j$$

for $k = 1, 2$ and real valued coefficients $(\beta_{k,j})_{k=1,2, j=1, \dots, d}$. Here each $Z_j \sim \mathcal{U}[-1, 1]$ independently and both ε_1 and ε_2 follow a $\mathcal{N}(0, 1)$ -distribution independently.

- (3) For data generating processes H_3 and A_3 we let $g_1 = g_2 = 1$ and

$$f_k(w_1, \dots, w_d) = \sum_{j=1}^d \beta_{1,k,j} w_j + \beta_{2,k,j} w_j^2$$

for $k = 1, 2$ and real valued coefficients $(\beta_{\ell,k,j})_{\ell=1,2,k=1,2,j=1,\dots,d}$. Here each $Z_j \sim \mathcal{U}[-1, 1]$ independently and both ε_1 and ε_2 follow a $\mathcal{N}(0, 1)$ -distribution independently.

- (4) For data generating processes H_4 and A_4 we let $f_1 = f_2 = 0$ and

$$g_k(w_1, \dots, w_d) = \sum_{j=1}^d \beta_{1,k,j} w_j + \beta_{2,k,j} w_j^2$$

for $k = 1, 2$ for real valued coefficients $(\beta_{\ell,k,j})_{\ell=1,2,k=1,2,j=1,\dots,d}$. Here each $Z_j \sim \mathcal{U}[-1, 1]$ independently and both ε_1 and ε_2 follow a $\mathcal{N}(0, 1)$ -distribution independently.

Each time we simulate from data generating processes H_1, \dots, H_4 we first draw the coefficients of the functions f_k, g_k from a $\mathcal{N}(0, 1)$ -distribution in order to make the results independent of a certain combination of parameters. When we simulate from the data generating processes A_1, A_2 and A_3 we first draw the coefficients of f_k, g_k to be either -1 or 1 with equal probability in order to fix the signal to noise ratio between the predictors and responses. When simulating from A_4 we simulate the coefficients of g_k to be either -5 or 5 , because the conditional dependence lies in the variance for A_4 , and a stronger signal is needed to compare the power of the tests using the same samples sizes as for A_1, A_2 and A_3 .

The data generating processes H_2, H_3, H_4, A_2, A_3 and A_4 can be shown to satisfy Assumption 3 that is needed for Corollary 9, since they are linear (in the parameters) location-scale models with bounded covariates (Belloni and Chernozhukov, 2011, Section 2.5). The processes H_1 and A_1 are not of this form, since g_1 and g_2 are nonlinear in the parameters. However, we include these in the simulation study to test the robustness of the test.

5.3 Level and power of partial copula test

In this section we examine the level and power properties of the test $\hat{\Psi}_n$. We examine the performance of the test on data generating processes H_1 and A_1 for dimensions $d \in \{1, 5, 10\}$ of Z . The test is performed as described in Section 4.7. As the quantile regression model we use an additive model with a B-spline basis of each variable with 5 degrees of freedom, and we try $q \in \{1, \dots, 5\}$. The result of the simulations can be seen in Figure 2. We observe that for $d = 1$ all five tests obtain level asymptotically under H_1 , while for higher dimension $d \in \{5, 10\}$ the test with $q = 4$ has minor problems holding level. We also see that the p -values for all five tests tend to zero as the sample size increases under A_1 . The convergence rate of the p -value depends on the dimension d such that a higher dimension gives a slower convergence rate. In conclusion we observe that our test holds level under a complicated data generating distribution (H_1), where there is a nonlinear conditional mean and variance dependence and skewed error distributions with super-Gaussian tails. Moreover, the test has power against the alternative of conditional dependence (A_1), however, for $d = 1$ we

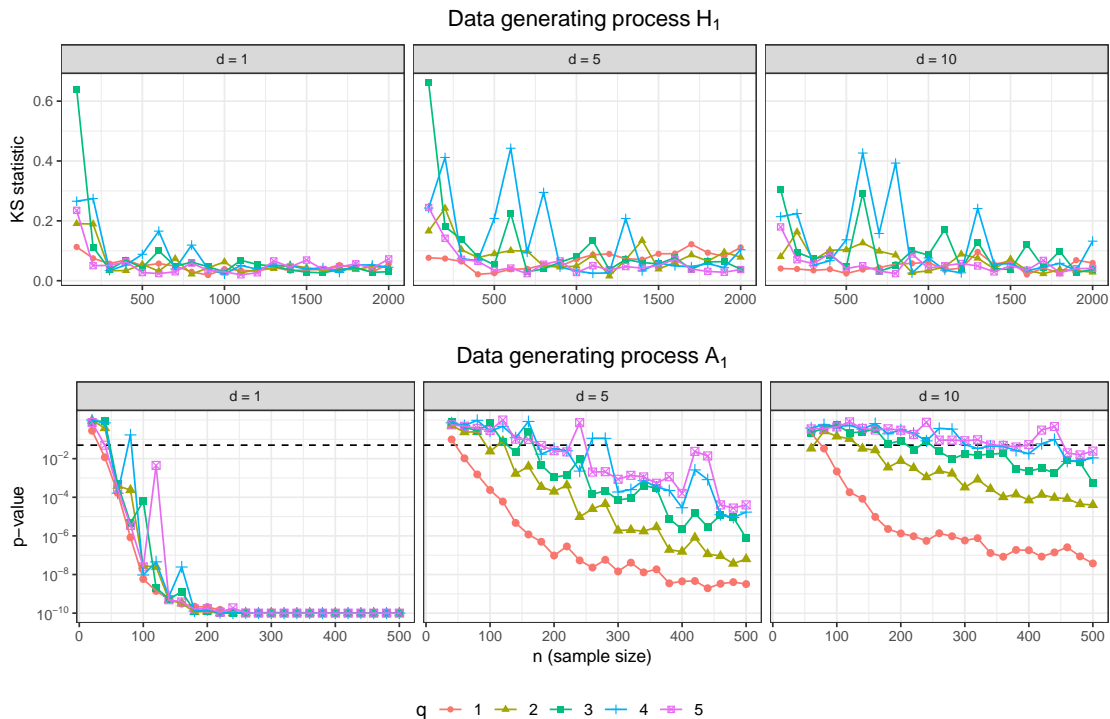


Figure 2: Top: KS statistic for equality with a $\mathcal{U}[0, 1]$ -distribution of the p-values of the two tests computed from 500 simulations from H_1 for each combination of n and d . Bottom: Average p-values of the two tests separately computed over 200 simulations from A_1 for each combination of n and d . Dashed line indicates the common significance level 0.05. For visual purposes all p -values have been truncated at 10^{-10} .

see that $q = 5$ gives the best power, while $q = 1$ gives the best power for $d \in \{5, 10\}$. The testing procedure also displays robustness to the fact that the quantile regression models are misspecified.

5.4 Comparison with other tests

We now compare the partial copula based test $\hat{\Psi}_n$ with other nonparametric tests. We will compare with a residual based method, since this is another class of conditional independence test based on nonparametric regression. In order to describe this test we let

$$R_{1,i} = X_i - \hat{f}(Z_i) \quad \text{and} \quad R_{2,i} = Y_i - \hat{g}(Z_i)$$

for $i = 1, \dots, n$ be the residuals obtained when performing conditional mean regression \hat{f} of $f(z) = E(X | Z = z)$ and \hat{g} of $g(z) = E(Y | Z = z)$ obtained from a sample $(X_i, Y_i, Z_i)_{i=1}^n$. We compare the following conditional independence tests:

- **GCM:** The Generalised Covariance Measure which tests for vanishing correlation between the residuals R_1 and R_2 given as above (Shah and Peters, 2020).
- **NPN correlation:** Testing for vanishing partial correlation in a nonparanormal distribution (Harris and Drton, 2013). This is a generalization of the partial correlation, which assumes a Gaussian dependence structure, but allows for arbitrary marginal distributions.
- **PC:** Our partial copula based test $\hat{\Psi}_n$ for $q \in \{1, 3, 5\}$ as described in Section 4.7.

We consider the behavior of the tests under H_2, A_2, H_3, A_3, H_4 and A_4 . For fairness of comparison we choose our quantile and mean regression models to be the correct model class such that the tests perform at their oracle level, e.g., for H_3 we fit additive models with polynomial basis of degree 2. We fix the dimension d of Z to be 3 in all simulations for simplicity. The results of the simulations can be seen in Figure 3.

Under H_2 all five tests hold level, and we see that the NPN test has greatest power against A_2 followed by the GCM and $\hat{\Psi}_n$ with $q = 1$, while $\hat{\Psi}_n$ with $q \in \{3, 5\}$ does not have much power against A_2 . In order to intuitively understand the effect of q see Figure 4. We see that in the estimated partial copula the dependence is captured by the overall correlation, while dividing $[\tau_{\min}, \tau_{\max}] \times [\tau_{\min}, \tau_{\max}]$ into subregions does not reveal finer dependence structure. Hence $q = 1$ is suitable to detect the dependence for A_1 .

Under H_3 both the GCM test and $\hat{\Psi}_n$ with $q \in \{1, 3, 5\}$ hold level, but the NPN test does not hold level under H_3 , which is due to the nonlinear response-predictor relationship. However, since both the GCM and $\hat{\Psi}_n$ test takes the nonlinearity into account, they can effectively filter away the Z -dependence. The NPN test has greatest power against the alternative A_3 following by $\hat{\Psi}_n$ with $q = 1$ and the GCM test. In Figure 4 we again see that the dependence in the estimated partial copula is described by the overall correlation, while dividing into subregions results a generalized correlation with elements that are close to zero, i.e., here $q = 1$ is suitable for capturing the dependence.

Under H_4 , all test hold level. Note that the NPN test holds level even though there is a nonlinear conditional variance relation, since this is still a nonparanormal distribution. We also see that neither the GCM test nor the NPN test has power against A_4 , while $\hat{\Psi}_n$ has some power against A_4 with the greatest power for $q = 3$. In Figure 4 we see that there is a clear dependence in the estimated partial copula, but that the overall correlation is close to zero. However, when dividing into subregions the generalized correlation is able to detect the dependencies in the tails of the distributions.

5.5 Power under local alternatives

Though GCM did not have power against the specific alternative A_4 , it maintains level and it has power against a broad class of alternatives. To understand better when $\hat{\Psi}_n$ can be expected to have greater power than GCM, we consider a simulation, which is a small variation of the simulations presented in Section 5.2.

The dimension is fixed as $d = 1$, $Z \sim \mathcal{U}([0, 1])$ is uniformly distributed on $[0, 1]$, $\varepsilon_1, \varepsilon_2$ and W are independent and $\mathcal{N}(0, 1)$ -distributed, and

$$X = (\beta Z^2 + 1)\varepsilon_1 + \gamma W \quad \text{and} \quad Y = (\beta Z^2 + 1)\varepsilon_2 + \gamma W \quad (\text{A})$$

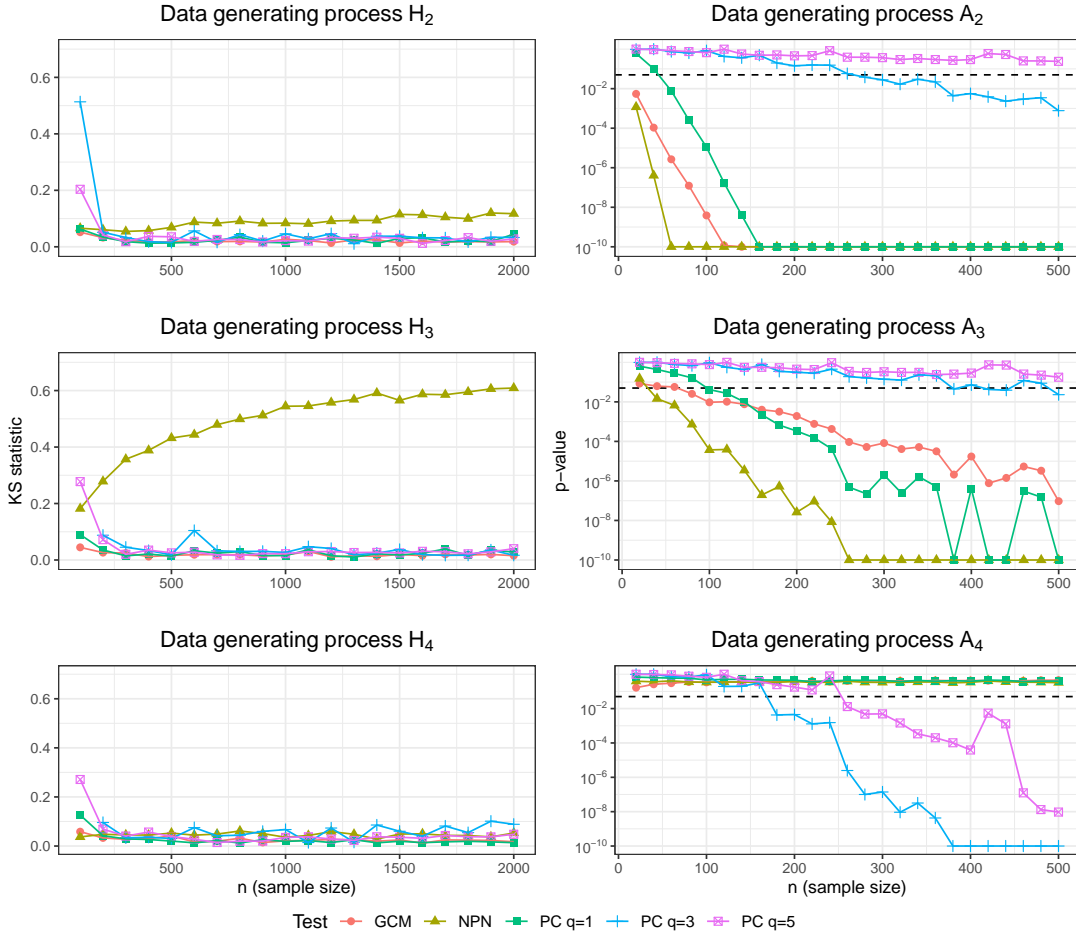


Figure 3: Left column: KS statistic for equality with a $\mathcal{U}[0, 1]$ -distribution of the p-values of the five tests computed from 500 simulations from H_2 , H_3 and H_4 , respectively, for each sample size n . Right column: Average p-values of the five tests separately computed over 200 simulations from A_2 , A_3 and A_4 , respectively, for each sample size n . For all simulations the dimension is fixed at $d = 3$. Dashed line indicates the common significance level 0.05. For visual purposes all p -values have been truncated at 10^{-10} .

for parameters $\beta, \gamma \in \mathbb{R}$. Conditionally on Z , the distribution of (X, Y) is a bivariate Gaussian distribution, and X and Y are conditionally independent if and only if $\gamma = 0$. We examine level and power by simulating 500 data sets for sample sizes $n \in \{100, 400, 1600\}$ and all combinations of parameters $\beta \in \{0, 1, 5, 10, 15, 20\}$, and local alternatives

$$\gamma^2 = \frac{\gamma_0^2}{\sqrt{n}}$$

for $\gamma_0^2 \in \{0, 50, 100, 150\}$. Note that $f(z) = E(X | Z = z) = 0$ and $g(z) = E(Y | Z = z) = 0$, which is exploited for GCM instead of estimating f and g . This should only increase the

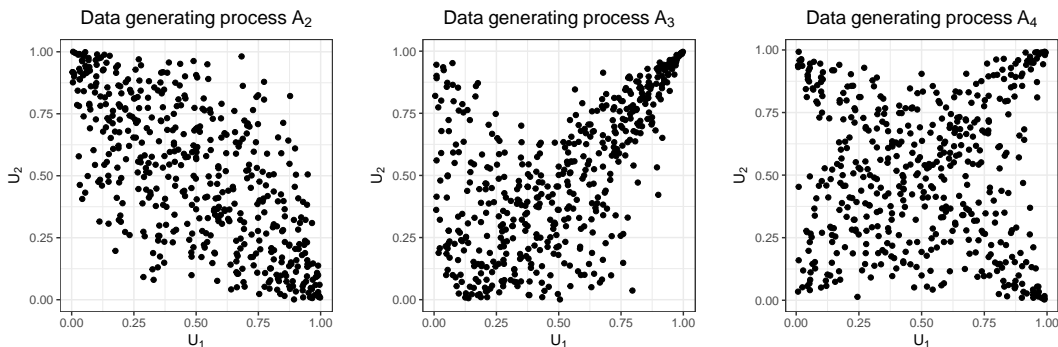


Figure 4: Estimated partial copula $(\hat{U}_{1,i}, \hat{U}_{2,i})_{i=1}^n$ from one realization from each of the data generating processes A_2, A_3 and A_4 for $d = 3$ and $n = 500$.

power of GCM relative to fitting any model of the conditional expectations. We perform the test $\hat{\Psi}_n$ as described in Section 4.7 using $q = 1$, and the quantile regression model is fitted using a polynomial basis of degree 2.

Figure 5 shows the results of the simulation. Both GCM and $\hat{\Psi}_n$ maintain level for $\gamma_0^2 = 0$. $\hat{\Psi}_n$ has comparable or superior power relative to GCM in all other cases. Both tests have decreasing power as a function of β , but $\hat{\Psi}_n$ maintains power even for large values of β , where GCM has almost no power. The power of $\hat{\Psi}_n$ against the local alternatives increases with the sample size, which shows how the increased precision for larger samples of the quantile regression based distribution functions improves power. We do not see the same for GCM, partly because no mean value model is fitted.

As β quantifies the conditional variance heterogeneity of X and Y given Z , we conclude that though GCM remains a valid test under conditional variance heterogeneity, its test statistic does not adequately account for the heterogeneity, and GCM has inferior power under local alternatives when compared to $\hat{\Psi}_n$.

6. Discussion

The first main contribution of this paper is an estimator of conditional distribution functions $\hat{F}^{(m,n)}$ based on quantile regression. We have shown that the estimator is pointwise (uniformly) consistent over a set of distributions $\mathcal{P}_0 \subset \mathcal{P}$ given that the quantile regression procedure is pointwise (uniformly) consistent over \mathcal{P}_0 . Moreover, we showed that the convergence rate of the quantile regression procedure can be transferred directly to the estimator $\hat{F}^{(m,n)}$.

The second main contribution of this paper is an analysis of a nonparametric test for conditional independence based on the partial copula construction. We introduced a class of tests given in terms of a generalized correlation dependence measure ρ with the leading example being a trimmed version of the Spearman correlation. We showed that the test achieves asymptotic pointwise (uniform) level and power over \mathcal{P}_0 given that the conditional distribution function estimators are pointwise (uniformly) consistent over \mathcal{P}_0 with rate functions g_P and h_P satisfying $\sqrt{ng_P(n)h_P(n)} \rightarrow 0$. The partial copula has previously

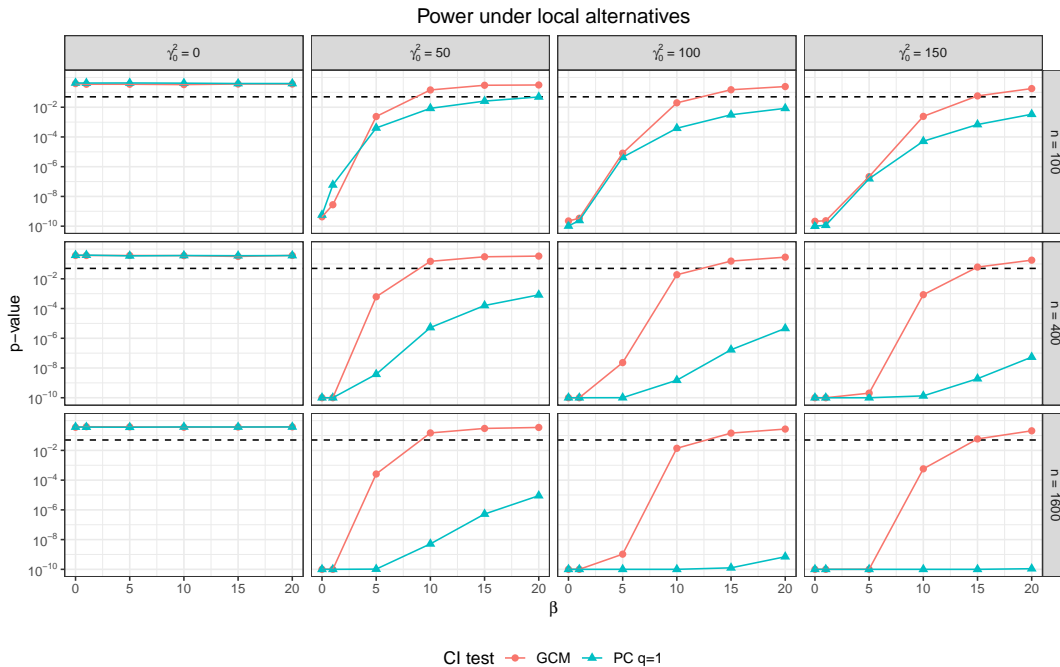


Figure 5: Average p-values of the three tests separately computed over 500 simulations for different values of the parameters β and γ_0 and different sample sizes n . The value $\gamma_0^2 = 0$ is equivalent to conditional independence. Other values correspond to the local alternatives $\gamma^2 = \gamma_0^2/\sqrt{n}$. The values of β determine the amount of heterogeneity of the mean regression residual variances with $\beta = 0$ meaning constant residual variance and $\beta = 20$ meaning substantial heterogeneity. Dashed line indicates the common significance level 0.05. For visual purposes all p -values have been truncated at 10^{-10} .

been considered for conditional independence testing in the literature, however, to the best of our knowledge, the results presented here are the first to explicitly connect the consistency requirements of the conditional distribution function estimators to level and power properties of the test.

Lastly, we established through a simulation study that the proposed test is sound under complicated data generating distributions, and that it has power comparable to or even better than other state-of-the-art nonparametric conditional independence tests. In particular, we demonstrated that our test has superior power against alternatives with variance heterogeneity between X and Y given Z when compared to conditional independence tests based on conventional residuals. We note that due to Daudin’s lemma, tests based on conventional residuals can obtain power against any alternative if suitable transformations of X and Y are considered. In particular, if X^2 and Y^2 were used in our simulation study, GCM would have power against A_4 . We tested the use of GCM in combination with X^2 and Y^2 in all our simulations (data not shown), and though it had some power against A_4 , it was comparable to or inferior to just using GCM in all other simulations. Thus to

obtain good power properties, the specific choice of transformation appears important and to depend on the data generating distribution.

An important point about the test is the rate requirement $\sqrt{ng_P(n)}h_P(n) \rightarrow 0$ needed to achieve asymptotic level. The product structure means that the test is sound under quantile regression models with slower consistency rates than the usual parametric $n^{-1/2}$ -rate. This opens up the methodology to nonparametric machine learning models. An interesting direction of research would be to empirically assess the performance of the test using machine learning inspired quantile regression models, such as deep neural networks, where explicit consistency rates are not available. We hypothesize that the method will still perform well in these scenarios due to the weak consistency requirement.

In this paper we have considered univariate X and Y . A possible extension of the test is to allow $X \in [0, 1]^{r_1}$ and $Y \in [0, 1]^{r_2}$ with $r_1, r_2 \geq 1$, and then consider the nonparametric residual $U_1 \in [0, 1]^{r_1}$ of X given Z by performing coordinatewise probability integral transformations $U_{1,k} = F_{X_k|Z}(X_k | Z)$ for $k = 1, \dots, r_1$, and similarly for constructing the nonparametric residual $U_2 \in [0, 1]^{r_2}$ of Y given Z . Conditional independence $X \perp\!\!\!\perp Y | Z$ then implies pairwise independence of $U_{1,k}$ and $U_{2,l}$ for each $k = 1, \dots, r_1$ and $l = 1, \dots, r_2$. Combining our proposed test statistics for each such pair yields an $r_1 r_2 q^2$ -dimensional test statistic, whose distribution under the hypothesis of conditional independence will be asymptotically Gaussian with mean 0. Its covariance matrix will only be partially known, though, due to the potential dependence between the pairs, but the unknown part could be estimated from the estimated nonparametric residuals. The multivariate statistic could be aggregated into a univariate test statistic in various ways, e.g. by a quadratic transformation as in (9), or by the maximum of the absolute values of its coordinates. In the low-dimensional case for fixed r_1 and r_2 our results would carry over immediately, and we expect that using the maximum could lead to high-dimensional results similar to Theorem 9 by Shah and Peters (2020).

A key property of the partial copula is that the nonparametric residuals U_1 and U_2 are independent under conditional independence and not only uncorrelated, which is the case for conventional residuals in additive noise models. Therefore, an important question is whether asymptotic level and power guaranties can be proven, when combining the partial copula with more general independence tests. In this paper we have focused on dependence measures of the form $\rho = E_P(\varphi(U_1)\varphi(U_2)^T)$ and tests based on

$$\hat{\rho}_n = \frac{1}{n} \sum_{i=1}^n \varphi(\hat{U}_{1,i})\varphi(\hat{U}_{2,i})^T$$

because it gives a flexible and general test for independence in the partial copula, it can be computed in linear time in the size of data, and most importantly its asymptotic theory is standard and easy to establish and apply. It also clearly illustrates the transfer of consistency of the conditional distribution function estimators to properties of the test. It is ongoing work to establish a parallel asymptotic theory for dependence measures of the form $\theta = E_P(h(U_1, U_2))$, where h is a kernel function, and whose estimators are U -statistics. This could potentially yield more powerfull tests against complicated alternatives of conditional dependence, but at the prize of increased computational complexity.

Acknowledgments

This work was supported by a research grant (13358) from VILLUM FONDEN.

Appendix A. Proofs

This appendix gives proofs of the main results of the paper. Throughout the proofs we will ignore the dependence of certain terms on the sample size to ease notation, e.g. we write $\hat{U}_{1,i}$ instead of $\hat{U}_{1,i}^{(n)}$ and $\hat{q}_{k,z}$ instead of $\hat{q}_{k,z}^{(n)}$.

A.1 Proof of Proposition 1

We need to bound the supremum

$$\|F - \tilde{F}^{(m)}\|_{\mathcal{T},\infty} = \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}|z)} |F(t|z) - \tilde{F}^{(m)}(t|z)|.$$

First we fix $z \in [0,1]^d$ and inspect the inner supremum. By construction we have

$$F(q_{k,z}|z) = \tilde{F}^{(m)}(q_{k,z}|z) = \tau_k$$

for $k = 1, \dots, m$. Furthermore, since both F and $\tilde{F}^{(m)}$ are continuous and increasing in $t \in [0,1]$ we have that

$$\sup_{t \in [q_{k,z}, q_{k+1,z}]} |F(t|z) - \tilde{F}^{(m)}(t|z)| \leq \tau_{k+1} - \tau_k$$

for each $k = 1, \dots, m-1$. Since $Q(\mathcal{T}|z) = [q_{\min,z}, q_{\max,z}] = \bigcup_{k=1}^{m-1} [q_{k,z}, q_{k+1,z}]$ we now have

$$\begin{aligned} \sup_{t \in Q(\mathcal{T}|z)} |F(t|z) - \tilde{F}^{(m)}(t|z)| &= \max_{k=1, \dots, m-1} \sup_{t \in [q_{k,z}, q_{k+1,z}]} |F(t|z) - \tilde{F}^{(m)}(t|z)| \\ &\leq \max_{k=1, \dots, m-1} (\tau_{k+1} - \tau_k) = \kappa_m. \end{aligned}$$

The result now follows from taking supremum over $z \in [0,1]^d$ as the right hand side of the inequality does not depend on z . \square

A.2 Proof of Proposition 4

We need to bound the supremum

$$\|\tilde{F}^{(m)} - \hat{F}^{(m,n)}\|_{\mathcal{T},\infty} = \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}|z)} |\tilde{F}^{(m)}(t|z) - \hat{F}^{(m,n)}(t|z)|.$$

Our proof strategy is the following. First we evaluate the inner supremum over $t \in Q(\mathcal{T}|z)$ analytically to obtain a bound in terms of the quantile regression prediction error. Then we will evaluate the outer supremum over $z \in [0,1]^d$ and use the assumed consistency from Assumption 1. First define the two quantities

$$A(m, n, z) := \kappa_m \cdot \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}^{(n)}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})}$$

and

$$B(m, n, z) := \kappa_m \cdot \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}^{(n)}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z}^{(n)} - \hat{q}_{k,z}^{(n)})}.$$

We then have the following key result regarding the inner supremum over $t \in Q(\mathcal{T} | z)$.

Proposition 24 *Let Assumption 1 (i) be satisfied. Then for all $P \in \mathcal{P}_0$ and $\varepsilon > 0$ there exists $N \geq 1$ such that for all $n \geq N$,*

$$\sup_{t \in Q(\mathcal{T} | z)} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \leq \max\{A(m, n, z), B(m, n, z)\}$$

for all $z \in [0, 1]^d$ and all grids $(\tau_k)_{k=1}^m$ in \mathcal{T} with probability at least $1 - \varepsilon$.

We need a number of auxilliary results before proving Proposition 24. We start by proving the following key lemma that reduces the number of distinct cases of relative positions of the true conditional quantiles $q_{k,z}$ and the estimated conditional quantiles $\hat{q}_{k,z}$.

Lemma 25 *Let Assumption 1 (i) be satisfied. Then for each $P \in \mathcal{P}_0$ and $\varepsilon > 0$ there exists $N \geq 1$ such that for all $n \geq N$ we have that $\hat{q}_{k,z} \in (q_{k-1,z}, q_{k+1,z})$ for each $k = 1, \dots, m$ and $z \in [0, 1]^d$ and for all grids $(\tau_k)_{k=1}^m$ in \mathcal{T} with probability at least $1 - \varepsilon$.*

Proof Fix a distribution $P \in \mathcal{P}_0$. Let \mathcal{G} be the set of all grids $(\tau_k)_{k=1}^m$ in \mathcal{T} . Then

$$\sup_{\mathcal{G}} \sup_{z \in [0, 1]^d} \max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}^{(n)}| \leq \sup_{z \in [0, 1]^d} \sup_{\tau \in \mathcal{T}} |Q(\tau | z) - \hat{Q}^{(n)}(\tau | z)| \xrightarrow{P} 0$$

under Assumption 1 (i). Since $q_{k,z} \in (q_{k-1,z}, q_{k+1,z})$ for each $k = 1, \dots, m$ and $z \in [0, 1]^d$ for all grids $(\tau_k)_{k=1}^m$ in \mathcal{T} the result follows. \blacksquare

Next we have some lemmas giving the supremum of certain functions over certain intervals that will be useful in the main proof.

Lemma 26 *Let $a \leq b < c \leq d$ and $f(t) = \frac{t-a}{c-a} - \frac{t-b}{d-b}$. Then $\sup_{t \in [b, c]} f(t) = \max\{\frac{b-a}{c-a}, \frac{d-c}{d-b}\}$.*

Proof Note that f is a linear function. Thus the supremum is obtained in one of the intervals endpoints, i.e., $\sup_{t \in [b, c]} f(t) = \max\{f(b), f(c)\}$. We see that

$$f(b) = \frac{b-a}{c-a} \quad \text{and} \quad f(c) = 1 - \frac{c-b}{d-b} = 1 - \frac{c-d+d-b}{d-b} = \frac{d-c}{d-b},$$

which shows the result. \blacksquare

Lemma 27 *Let $a < b \leq c < d$ and $f(t) = \alpha + \beta \cdot \frac{t-b}{d-b} - \alpha \cdot \frac{t-a}{c-a}$ where $\alpha, \beta > 0$. Then we have $\sup_{t \in [b, c]} f(t) = \max\{\alpha \cdot \frac{c-b}{c-a}, \beta \cdot \frac{c-b}{d-b}\}$.*

Proof The function f is a linear function, and hence the supremum is obtained in one of the interval endpoints. We see that

$$f(b) = \alpha - \alpha \cdot \frac{b-a}{c-a} = \alpha \cdot \frac{c-b}{c-a} \quad \text{and} \quad f(c) = \beta \cdot \frac{c-b}{d-b},$$

which shows the claim. ■

Lemma 28 *Let $a \leq b < c < d$ and $f(t) = |g(t)|$ where $g(t) = \frac{t-b}{c-b} - \frac{t-a}{d-a}$. Then we have that $\sup_{t \in [b,c]} f(t) = \max\{\frac{b-a}{d-a}, \frac{d-c}{d-a}\}$.*

Proof Note that $f(t)$ is a convex function. Therefore the supremum of $f(t)$ is obtained in one of the interval endpoints. We see that

$$f(b) = \left| \frac{b-b}{c-b} - \frac{b-a}{d-a} \right| = \left| -\frac{b-a}{d-a} \right| = \frac{b-a}{d-a}$$

and

$$f(c) = \left| \frac{c-b}{c-b} - \frac{c-a}{d-a} \right| = \left| 1 - \frac{c-a}{d-a} \right| = \left| 1 - \frac{c-d+d-a}{d-a} \right| = \frac{d-c}{d-a}$$

which was what we wanted. ■

We are now ready to show Proposition 24.

Proof [Proof (of Proposition 24)]

We will compute the supremum over $t \in Q(\mathcal{T} | z) = [q_{\min,z}, q_{\max,z}]$ as the maximum of the suprema over the intervals $[q_{k,z}, q_{k+1,z}]$ for $k = 1, \dots, m-1$, i.e.,

$$\sup_{t \in Q(\mathcal{T}|z)} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| = \max_{k=1, \dots, m-1} \sup_{t \in [q_{k,z}, q_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)|.$$

This is useful since on each interval of the form $[q_{k,z}, q_{k+1,z}]$ we have that $\tilde{F}^{(m)}(\cdot | z)$ is a linear function, while $\hat{F}^{(m,n)}(\cdot | z)$ is a piecewise linear function.

First fix a distribution $P \in \mathcal{P}_0$ and $\varepsilon > 0$. Using Lemma 25 we choose $N \geq 1$ such that $\hat{q}_{k,z} \in (q_{k-1,z}, q_{k+1,z})$ for $k = 1, \dots, m-1$ and $z \in [0, 1]^d$ for each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} with probability at least $1 - \varepsilon$. Now fix a $k = 1, \dots, m-1$ such that we will examine the supremum on $[q_{k,z}, q_{k+1,z}]$. The relative position of the true and estimated conditional quantiles can be divided into four cases:

- 1) $q_{k,z} \geq \hat{q}_{k,z}$ and $q_{k+1,z} \geq \hat{q}_{k+1,z}$.
- 2) $q_{k,z} \geq \hat{q}_{k,z}$ and $q_{k+1,z} < \hat{q}_{k+1,z}$.
- 3) $q_{k,z} < \hat{q}_{k,z}$ and $q_{k+1,z} \geq \hat{q}_{k+1,z}$.
- 4) $q_{k,z} < \hat{q}_{k,z}$ and $q_{k+1,z} < \hat{q}_{k+1,z}$.

We start with case 1). First we compute the supremum over $t \in [q_{k,z}, \hat{q}_{k+1,z}]$ and then over $t \in [\hat{q}_{k+1,z}, q_{k+1,z}]$. We have that

$$|\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| = (\tau_{k+1} - \tau_k) \left(\frac{t - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} - \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} \right)$$

for $t \in [q_{k,z}, \hat{q}_{k+1,z}]$. Hence we can compute the supremum as

$$\begin{aligned} & \sup_{t \in [q_{k,z}, \hat{q}_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_{k+1} - \tau_k) \max \left\{ \frac{q_{k,z} - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}}, \frac{q_{k+1,z} - \hat{q}_{k+1,z}}{q_{k+1,z} - q_{k,z}} \right\} \\ &\leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used Lemma 26. Now we see that

$$\begin{aligned} & |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_{k+1} - \tau_k) + (\tau_{k+2} - \tau_{k+1}) \frac{t - \hat{q}_{k+1,z}}{\hat{q}_{k+2,z} - \hat{q}_{k+1,z}} - (\tau_{k+1} - \tau_k) \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} \end{aligned}$$

for $t \in [\hat{q}_{k+1,z}, q_{k+1,z}]$. We compute the supremum to be

$$\begin{aligned} & \sup_{t \in [\hat{q}_{k+1,z}, q_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= \max \left\{ (\tau_{k+1} - \tau_k) \frac{\hat{q}_{k+1,z} - q_{k+1,z}}{q_{k+1,z} - q_{k,z}}, (\tau_{k+2} - \tau_{k+1}) \frac{\hat{q}_{k+1,z} - q_{k+1,z}}{\hat{q}_{k+2,z} - \hat{q}_{k+1,z}} \right\} \\ &\leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used Lemma 27. This covers case 1).

Now let us proceed to case 2). Here we can evaluate the supremum over $t \in [q_{k,z}, q_{k+1,z}]$ directly. We have that

$$|\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| = (\tau_{k+1} - \tau_k) \left| \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} - \frac{t - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} \right|.$$

The supremum can now be evaluated using Lemma 28 to be

$$\begin{aligned} & \sup_{t \in [q_{k,z}, q_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_{k+1} - \tau_k) \max \left\{ \frac{q_{k,z} - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}}, \frac{\hat{q}_{k+1,z} - q_{k+1,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} \right\} \\ &\leq \kappa_m \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}. \end{aligned}$$

In case 3) we need to divide into three cases, namely when $t \in [q_{k,z}, \hat{q}_{z,k}]$, $t \in [\hat{q}_{k,z}, \hat{q}_{k+1,z}]$ and $t \in [\hat{q}_{k+1,z}, q_{k+1,z}]$. In the first case we have

$$\begin{aligned} & |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_k - \tau_{k-1}) + (\tau_{k+1} - \tau_k) \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} - (\tau_k - \tau_{k-1}) \frac{t - \hat{q}_{k-1,z}}{\hat{q}_{k,z} - \hat{q}_{k-1,z}} \end{aligned}$$

for $t \in [q_{k,z}, \hat{q}_{k,z}]$. Therefore we have

$$\begin{aligned} & \sup_{t \in [q_{k,z}, \hat{q}_{z,k}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ & \leq \max \left\{ (\tau_k - \tau_{k-1}) \frac{\hat{q}_{k,z} - q_{k,z}}{\hat{q}_{k,z} - \hat{q}_{k-1,z}}, (\tau_{k+1} - \tau_k) \frac{\hat{q}_{k,z} - q_{k,z}}{q_{k+1,z} - q_{k,z}} \right\} \\ & \leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used Lemma 27. In the second case we have

$$|\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| = (\tau_{k+1} - \tau_k) \left| \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} - \frac{t - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} \right|,$$

for $t \in [\hat{q}_{k,z}, \hat{q}_{k+1,z}]$ and therefore we obtain

$$\begin{aligned} & \sup_{t \in [\hat{q}_{k,z}, \hat{q}_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ & \leq (\tau_{k+1} - \tau_k) \max \left\{ \frac{\hat{q}_{k,z} - q_{k,z}}{q_{k+1,z} - q_{k,z}}, \frac{q_{k+1,z} - \hat{q}_{k+1,z}}{q_{k+1,z} - q_{k,z}} \right\} \\ & \leq \kappa_m \cdot \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \end{aligned}$$

where we have used Lemma 28. In the third case we have

$$\begin{aligned} & |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_{k+1} - \tau_k) + (\tau_{k+2} - \tau_{k+1}) \frac{t - \hat{q}_{k+1,z}}{\hat{q}_{k+2,z} - \hat{q}_{k+1,z}} - (\tau_{k+1} - \tau_k) \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} \end{aligned}$$

for $t \in [\hat{q}_{k+1,z}, q_{k+1,z}]$. So we obtain

$$\begin{aligned} & \sup_{t \in [\hat{q}_{k+1,z}, q_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= \max \left\{ (\tau_{k+1} - \tau_k) \frac{\hat{q}_{k+1,z} - q_{k+1,z}}{q_{k+1,z} - q_{k,z}}, (\tau_{k+2} - \tau_{k+1}) \frac{\hat{q}_{k+1,z} - q_{k+1,z}}{\hat{q}_{k+2,z} - \hat{q}_{k+1,z}} \right\} \\ & \leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used Lemma 27.

Let us now examine case 4). Here we have the two sub cases $t \in [q_{k,z}, \hat{q}_{k,z}]$ and $t \in [\hat{q}_{k,z}, q_{k+1,z}]$. First we see that

$$\begin{aligned} & |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_k - \tau_{k-1}) + (\tau_{k+1} - \tau_k) \frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} - (\tau_k - \tau_{k-1}) \frac{t - \hat{q}_{k-1,z}}{\hat{q}_{k,z} - \hat{q}_{k-1,z}} \end{aligned}$$

for $t \in [q_{k,z}, \hat{q}_{k,z}]$. Thus we have

$$\begin{aligned} & \sup_{t \in [q_{k,z}, \hat{q}_{k,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ & \leq \max \left\{ (\tau_k - \tau_{k-1}) \frac{\hat{q}_{k,z} - q_{k,z}}{\hat{q}_{k,z} - \hat{q}_{k-1,z}}, (\tau_{k+1} - \tau_k) \frac{\hat{q}_{k,z} - q_{k,z}}{q_{k+1,z} - q_{k,z}} \right\} \\ & \leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used 27. Now in the second case we have

$$|\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| = (\tau_{k+1} - \tau_k) \left(\frac{t - q_{k,z}}{q_{k+1,z} - q_{k,z}} - \frac{t - \hat{q}_{k,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} \right)$$

for $t \in [\hat{q}_{k,z}, q_{k+1,z}]$. From this we get the supremum to be

$$\begin{aligned} & \sup_{t \in [\hat{q}_{k,z}, q_{k+1,z}]} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \\ &= (\tau_{k+1} - \tau_k) \max \left\{ \frac{\hat{q}_{k,z} - q_{k,z}}{q_{k+1,z} - q_{k,z}}, \frac{q_{k+1,z} - \hat{q}_{k+1,z}}{\hat{q}_{k+1,z} - \hat{q}_{k,z}} \right\} \\ & \leq \kappa_m \max \left\{ \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z})}, \frac{\max_{k=1, \dots, m} |q_{k,z} - \hat{q}_{k,z}|}{\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z})} \right\} \end{aligned}$$

where we have used Lemma 26. Taking maximum of all cases and sub cases yields the desired result. \blacksquare

We will now move on to tackling the problem of controlling the outer supremum over $z \in [0, 1]^d$. First we prove the following technical lemma that gives control over the denominators in $A(m, n, z)$ and $B(m, n, z)$.

Lemma 29 *Let Assumption 1 be satisfied. Let $\gamma_m := \min_{k=1, \dots, m-1} (\tau_{k+1} - \tau_k)$ denote the finest subinterval of the grid. Then for each $P \in \mathcal{P}_0$ we have*

$$\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z}) \geq \frac{\gamma_m}{C_P}$$

for almost all $z \in [0, 1]^d$ for each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Also for all $\varepsilon > 0$ there is $N \geq 1$ such that for all $n \geq N$ we have

$$\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z}) \geq \frac{\gamma_m}{3C_P}$$

for almost all $z \in [0, 1]^d$ for each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} with probability at least $1 - \varepsilon$.

Proof Fix a distribution $P \in \mathcal{P}_0$. We see that

$$\begin{aligned}\tau_{k+1} - \tau_k &= F(q_{k+1,z} | z) - F(q_{k,z} | z) \\ &= \int_{q_{k,z}}^{q_{k+1,z}} f(x | z) dx \leq C_P \cdot (q_{k+1,z} - q_{k,z})\end{aligned}$$

for each $k = 1, \dots, m-1$ and almost all $z \in [0, 1]^d$ for each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Here we have used Assumption 1 (ii). Rearranging and taking minimum, we have that

$$\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z}) \geq \min_{k=1, \dots, m-1} \frac{\tau_{k+1} - \tau_k}{C_P} = \frac{\gamma_m}{C_P}$$

for almost all $z \in [0, 1]^d$ and each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Now let $\varepsilon > 0$ be given. Choose $N \geq 1$ such that for all $n \geq N$ we have

$$P \left(\hat{q}_{k,z} \in \left(q_{k,z} - \frac{\gamma_m}{3C_P}, q_{k,z} + \frac{\gamma_m}{3C_P} \right) \right) \geq 1 - \varepsilon$$

for all $k = 1, \dots, m$ and all $z \in [0, 1]^d$ for each $(\tau_k)_{k=1}^m$ in \mathcal{T} , which is possible due to Assumption 1 (i). In this case

$$\hat{q}_{k,z} \leq q_{k,z} + \frac{\gamma_m}{3C_P} \quad \text{and} \quad \hat{q}_{k+1,z} \geq q_{k+1,z} - \frac{\gamma_m}{3C_P}$$

for all $k = 1, \dots, m-1$ and $z \in [0, 1]^d$ with probability at least $1 - \varepsilon$. Thus for $n \geq N$,

$$\begin{aligned}\min_{k=1, \dots, m-1} (\hat{q}_{k+1,z} - \hat{q}_{k,z}) &\geq \min_{k=1, \dots, m-1} \left(q_{k+1,z} - \frac{\gamma_m}{3C_P} - \left(q_{k,z} + \frac{\gamma_m}{3C_P} \right) \right) \\ &= \min_{k=1, \dots, m-1} \left(q_{k+1,z} - q_{k,z} - \frac{2\gamma_m}{3C_P} \right) \\ &\geq \frac{\gamma_m}{C_P} - \frac{2\gamma_m}{3C_P} = \frac{\gamma_m}{3C_P}\end{aligned}$$

for all $z \in [0, 1]^d$ and each grid $(\tau_k)_{k=1}^m$ in \mathcal{T} with probability at least $1 - \varepsilon$. ■

We are now ready to prove the main result.

Proof [Proof (of Proposition 4)]

Fix a distribution $P \in \mathcal{P}_0$. Let $\varepsilon \in (0, 1)$ be given. Firstly, we use Proposition 24 to choose $N_1 \geq 1$ such that the event

$$E_1 = \left(\sup_{t \in Q(\mathcal{T}|z)} |\tilde{F}^{(m)}(t | z) - \hat{F}^{(m,n)}(t | z)| \leq \max\{A(m, n, z), B(m, n, z)\} \right)$$

has probability at least $1 - \varepsilon/3$ for all $n \geq N_1$ and every grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Secondly, according to Lemma 29 we have that

$$\min_{k=1, \dots, m-1} (q_{k+1,z} - q_{k,z}) \geq \frac{\gamma_m}{C_P}$$

and we can choose $N_2 \geq 1$ such that the event

$$E_2 = \left(\min_{k=1, \dots, m-1} (\hat{q}_{k+1, z} - \hat{q}_{k, z}) \geq \frac{\gamma_m}{3C_P} \right)$$

has probability at least $1 - \varepsilon/3$ for all $n \geq N_2$ and every grid $(\tau_k)_{k=1}^m$ in \mathcal{T} . Thirdly, we can choose $N_3 \geq 1$ and $M'_P > 0$ such that the event

$$E_3 = \left(\frac{\sup_{z \in [0, 1]^d} \max_{k=1, \dots, m} |q_{k, z} - \hat{q}_{k, z}|}{g_P(n)} \leq M'_P \right)$$

has probability at least $1 - \varepsilon/3$ for all $n \geq N_3$ and every $(\tau_k)_{k=1}^m$ in \mathcal{T} using Assumption 1 (i). Now we note that on the event $E := E_1 \cap E_2 \cap E_3$ we have

$$\begin{aligned} \|\tilde{F}^{(m)} - \hat{F}^{(m, n)}\|_{\mathcal{T}, \infty} &\leq \sup_{z \in [0, 1]^d} \max\{A(m, n, z), B(m, n, z)\} \\ &= 3C_P \cdot \frac{\kappa_m}{\gamma_m} \sup_{z \in [0, 1]^d} \max_{k=1, \dots, m} |q_{k, z} - \hat{q}_{k, z}| \\ &\leq 3C_P \cdot M'_P \cdot g_P(n) \end{aligned}$$

with probability $P(E) \geq 1 - \varepsilon$ for all $n \geq N$ and every grid $(\tau_k)_{k=1}^m$ in \mathcal{T} where $N := \max\{N_1, N_2, N_3\}$. Here we have used that $\kappa_m/\gamma_m = 1$ due to the grids being equidistant. We can now set $M_P := 3C_P \cdot M'_P$ such that

$$P \left(\frac{\|\tilde{F}^{(m)} - \hat{F}^{(m, n)}\|_{\mathcal{T}, \infty}}{g_P(n)} > M_P \right) < \varepsilon$$

whenever $n \geq N$. This shows that $\|\tilde{F}^{(m)} - \hat{F}^{(m, n)}\|_{\mathcal{T}, \infty} \in \mathcal{O}_P(g_P(n))$ for every equidistant grid $(\tau_k)_{k=1}^m$ in \mathcal{T} as wanted. \blacksquare

A.3 Proof of Theorem 5

According to Corollary 3 we have

$$\|F - \hat{F}^{(m_n, n)}\|_{\infty} \leq \kappa_{m_n} + \|\tilde{F}^{(m_n)} - \hat{F}^{(m_n, n)}\|_{\infty}.$$

Here $\|\tilde{F}^{(m_n)} - \hat{F}^{(m_n, n)}\|_{\infty} \in \mathcal{O}_P(g_P(n))$ for each equidistant grid $(\tau_k)_{k=1}^{m_n}$ in \mathcal{T} due to Proposition 4. Since we have assumed that $\kappa_{m_n} \in o(g_P(n))$ we have the result. \square

A.4 Proof of Proposition 6

The proof follows immediately from the proof of Proposition 4 and the stronger Assumption 2 in the following way. Note that the statement of Lemma 25 holds uniformly over $P \in \mathcal{P}_0$ under Assumption 2 (i). Therefore Proposition 24 also holds uniformly over $P \in \mathcal{P}_0$. Furthermore, the result of Lemma 29 also holds uniformly in $P \in \mathcal{P}_0$ under Assumption 2. Therefore the probability of the events E_1, E_2 and E_3 can be controlled uniformly over $P \in \mathcal{P}_0$ from which the result follows. \square

A.5 Proof of Theorem 7

The corollary follows from Proposition 6 using the same argument as in the proof of Theorem 5. \square

A.6 Proof of Corollary 9

Using Theorem 8 we have that

$$\begin{aligned} \sup_{z \in [0,1]^d} \sup_{\tau \in \mathcal{Q}} |Q(\tau | z) - \hat{Q}(\tau | z)| &= \sup_{z \in [0,1]^d} \sup_{\tau \in \mathcal{Q}} |h(z)^T (\beta_\tau - \hat{\beta}_\tau)| \\ &\leq \sup_{z \in [0,1]^d} \|h(z)\|_2 \sup_{\tau \in \mathcal{Q}} \|\beta_\tau - \hat{\beta}_\tau\|_2 \\ &\in \mathcal{O}_P \left(\sqrt{\frac{s_n \log(p \vee n)}{n}} \right) \end{aligned}$$

since $\sup_{z \in [0,1]^d} \|h(z)\|_2 < \infty$ because $[0, 1]^d$ is compact and h is continuous. \square

A.7 Proof of Proposition 11

Assume that $X \perp\!\!\!\perp Y | Z$. Then it also holds that $(X, Z) \perp\!\!\!\perp (Y, Z) \perp\!\!\!\perp Z$ and thus $U_1 \perp\!\!\!\perp U_2 | Z$. Letting f denote a generic density function, we now have that

$$\begin{aligned} f(u_1, u_2) &= \int f(u_1, u_2 | z) f(z) dz = \int f(u_1 | z) f(u_2 | z) f(z) dz \\ &= \int f(u_1) f(u_2) f(z) dz = f(u_1) f(u_2) \end{aligned}$$

for all $u_1, u_2 \in [0, 1]$, where we have used Proposition 10. \square

A.8 Proof of Theorem 14

Before proving the theorem, we will supply a lemma that will aid us during the proof.

Lemma 30 *Let $\hat{F}_{X|Z}^{(n)}$ and $\hat{F}_{Y|Z}^{(n)}$ satisfy Assumption 5. Then*

$$\|\varphi_k \circ F_{X|Z} - \varphi_k \circ \hat{F}_{X|Z}^{(n)}\|_\infty \in \mathcal{O}_P(g_P(n)) \quad \text{and} \quad \|\varphi_k \circ F_{Y|Z} - \varphi_k \circ \hat{F}_{Y|Z}^{(n)}\|_\infty \in \mathcal{O}_P(h_P(n))$$

for each $k = 1, \dots, q$ given that φ satisfies Assumption 4.

Proof We only show the first statement. Fix $k = 1, \dots, q$. We need to control the supremum

$$\sup_{z \in [0,1]^d} \sup_{t \in [0,1]} |\varphi_k(F_{X|Z}(t | z)) - \varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))|.$$

We will divide the supremum over $t \in [0, 1]$ into two cases. Namely, when $t \in Q(\mathcal{T} | z) = [q_{\min, z}, q_{\max, z}]$ and when $t \in Q(\mathcal{T}^c | z) = [q_{\min, z}, q_{\max, z}]^c$. First we see that

$$\begin{aligned} \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T} | z)} |\varphi_k(F_{X|Z}(t | z)) - \varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))| \\ \leq L_k \cdot \|F_{X|Z} - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}, \infty} \in \mathcal{O}_P(g_P(n)) \end{aligned}$$

where L_k is the Lipschitz constant of φ_k under Assumption 4 (ii). Here we have used the consistency in Assumption 5 (i). Next we examine the supremum over $t \in Q(\mathcal{T}^c | z)$. First note that $F_{X|Z}(t | z) \in [\tau_{\min}, \tau_{\max}]^c$ whenever $t \in Q(\mathcal{T}^c | z)$. Also recall that the support of φ_k is $\mathcal{T}_k \subset \mathcal{T} = [\tau_{\min}, \tau_{\max}]$. Therefore $\varphi_k(F_{X|Z}(t | z)) = 0$ for $t \in Q(\mathcal{T}^c | z)$. Hence we have

$$\sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}^c | z)} |\varphi_k(F_{X|Z}(t | z)) - \varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))| = \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}^c | z)} |\varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))|.$$

By Assumption 5 (i) we know that

$$\hat{F}_{X|Z}^{(n)}(q_{\min,z} | z) \xrightarrow{P} \tau_{\min} \quad \text{and} \quad \hat{F}_{X|Z}^{(n)}(q_{\max,z} | z) \xrightarrow{P} \tau_{\max}$$

for all $z \in [0,1]^d$. Since $\hat{F}_{X|Z}^{(n)}(\cdot | z)$ is increasing we thus know that the limit $\xi(t, z)$ from Assumption 5 (ii) must satisfy $\xi(t, z) \in [\tau_{\min}, \tau_{\max}]^c$ for $t \in Q(\mathcal{T}^c | z)$ and $z \in [0,1]^d$. Again, since the support of φ_k is $\mathcal{T}_k \subset \mathcal{T} = [\tau_{\min}, \tau_{\max}]$ we have that $\varphi_k(\xi(t, z)) = 0$ when $t \in Q(\mathcal{T}^c | z)$ and $z \in [0,1]^d$. Therefore we have that

$$\begin{aligned} \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}^c | z)} |\varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))| &= \sup_{z \in [0,1]^d} \sup_{t \in Q(\mathcal{T}^c | z)} |\varphi_k(\xi(t, z)) - \varphi_k(\hat{F}_{X|Z}^{(n)}(t | z))| \\ &\leq L_k \cdot \|\xi - \hat{F}_{X|Z}^{(n)}\|_{\mathcal{T}^c, \infty} \in \mathcal{O}_P(g_P(n)), \end{aligned}$$

where we have used Assumption 5 (ii). Putting the two cases together we have that

$$\|\varphi_k \circ F_{X|Z} - \varphi_k \circ \hat{F}_{X|Z}^{(n)}\|_{\infty} \in \mathcal{O}_P(g_P(n))$$

which was what we wanted. ■

We can now prove the main theorem.

Proof [Proof (of Theorem 14)] Fix a distribution $P \in \mathcal{H}_0$. The key to proving the theorem is the decomposition

$$\hat{\rho}_n = \alpha_n + \beta_n + \gamma_n + \delta_n$$

where $\alpha_n, \beta_n, \gamma_n$ and δ_n are given by

$$\begin{aligned} \alpha_n &= \frac{1}{n} \sum_{i=1}^n \varphi(U_{1,i}) \varphi(U_{2,i})^T, \\ \beta_n &= \frac{1}{n} \sum_{i=1}^n \left(\varphi(\hat{U}_{1,i}) - \varphi(U_{1,i}) \right) \left(\varphi(\hat{U}_{2,i}) - \varphi(U_{2,i}) \right)^T \\ \gamma_n &= \frac{1}{n} \sum_{i=1}^n \varphi(U_{1,i}) \left(\varphi(\hat{U}_{2,i}) - \varphi(U_{2,i}) \right)^T, \\ \delta_n &= \frac{1}{n} \sum_{i=1}^n \left(\varphi(\hat{U}_{1,i}) - \varphi(U_{1,i}) \right) \varphi(U_{2,i})^T \end{aligned}$$

The term α_n will be driving the asymptotics of the test statistics, while β_n, γ_n and δ_n are error terms that we wish to show converge to zero sufficiently fast.

Let us start by examining α_n . Under Assumption 4 (iii) we see that

$$E_P(\varphi(U_{1,i})\varphi(U_{2,i})^T) = E_P(\varphi(U_{1,i}))E_P(\varphi(U_{2,i}))^T = 0$$

because $P \in \mathcal{H}_0$ and furthermore we see that

$$\begin{aligned} \text{Cov}_P(\varphi_k(U_{1,i})\varphi_\ell(U_{2,i}), \varphi_s(U_{1,i})\varphi_t(U_{2,i})) &= E_P(\varphi_k(U_{1,i})\varphi_\ell(U_{2,i})\varphi_s(U_{1,i})\varphi_t(U_{2,i})) \\ &= E_P(\varphi_k(U_{1,i})\varphi_s(U_{1,i}))E_P(\varphi_\ell(U_{2,i})\varphi_t(U_{2,i})) \\ &= \int_0^1 \varphi_k(u)\varphi_s(u)du \int_0^1 \varphi_\ell(u)\varphi_t(u)du \\ &= \Sigma_{ks}\Sigma_{\ell t} = (\Sigma \otimes \Sigma)_{k\ell, st} \end{aligned}$$

for $k, \ell, s, t = 1, \dots, q$. Observe that $\Sigma_{k,k} = 1$. Since α_n is the average of i.i.d. terms with zero mean and covariance $\Sigma \otimes \Sigma$, the central limit theorem states that

$$\sqrt{n}\alpha_n \Rightarrow_P \mathcal{N}(0, \Sigma \otimes \Sigma)$$

for each $P \in \mathcal{H}_0$.

Now let us examine the term $\sqrt{n}\beta_n$. Fix $k, \ell = 1, \dots, q$. Then we have

$$\begin{aligned} |\sqrt{n}\beta_{k\ell, n}| &\leq \frac{1}{\sqrt{n}} \sum_{i=1}^n \left| \varphi_k(\hat{U}_{1,i}) - \varphi_k(U_{1,i}) \right| \cdot \left| \varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right| \\ &\leq \frac{n}{\sqrt{n}} \|\varphi_k \circ \hat{F}_{X|Z}^{(n)} - \varphi_k \circ F_{X|Z}\|_\infty \cdot \|\varphi_\ell \circ \hat{F}_{Y|Z}^{(n)} - \varphi_\ell \circ F_{Y|Z}\|_\infty \\ &\in \mathcal{O}_P(\sqrt{n}g_P(n)h_P(n)) \end{aligned}$$

where we have used Lemma 30, which is valid due to Assumption 5. Since we have assumed that the rate functions satisfy $\sqrt{n}g_P(n)h_P(n) \rightarrow 0$ we can conclude that $|\sqrt{n}\beta_{k\ell, n}| \rightarrow_P 0$ for each $k, \ell = 1, \dots, q$. Hence $\sqrt{n}\beta_n \rightarrow_P 0$.

Now we turn to the cross terms γ_n and δ_n . The two terms are dealt with analogously, so we only examine γ_n . Fix $k, \ell = 1, \dots, q$ and consider writing

$$\gamma_{k\ell, n} = \frac{1}{n} \sum_{i=1}^n C_i \quad \text{where} \quad C_i = \varphi_k(U_{1,i}) \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right).$$

We will compute the mean and variance of $\sqrt{n}\gamma_{k\ell, n}$ conditionally on $(Y_j, Z_j)_{j=1}^n$ in order to use Chebyshev's inequality to show that it converges to zero in probability. Observe that

$$\begin{aligned} E_P(C_i | (Y_j, Z_j)_{j=1}^n) &= E_P \left(\varphi_k(U_{1,i}) \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right) | (Y_j, Z_j)_{j=1}^n \right) \\ &= \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right) E_P \left(\varphi_k(U_{1,i}) | (Y_j, Z_j)_{j=1}^n \right) \quad \text{a.s.} \end{aligned}$$

Here we have exploited that $\varphi_\ell(U_{2,i})$ and $\varphi_\ell(\hat{U}_{2,i}) = \varphi_\ell(\hat{F}_{Y|Z}(Y_i | Z_i))$ are measurable functions of $(Y_j, Z_j)_{j=1}^n$. Now since $P \in \mathcal{H}_0$ we have $\varphi_k(U_{1,i}) \perp\!\!\!\perp Y_i | Z_i$ and $\varphi_k(U_{1,i}) \perp\!\!\!\perp Z_i$ due to Proposition 10. Therefore

$$E_P \left(\varphi_k(U_{1,i}) | (Y_j, Z_j)_{j=1}^n \right) = E_P(\varphi_k(U_{1,i})) = 0 \quad \text{a.s.}$$

where we have used Assumption 4 (iii). Hence $E_P(C_i \mid (Y_j, Z_j)_{j=1}^n) = 0$ a.s. From the tower property we also obtain that $E_P(C_i) = 0$ and therefore $\sqrt{n}\gamma_{kl,n}$ has mean zero. Let us turn to the conditional variance. Conditionally on $(Y_j, Z_j)_{j=1}^n$ the terms $(C_i)_{i=1}^n$ are i.i.d. because $\varphi_\ell \circ \hat{F}_{Y|Z}$ is $(Y_j, Z_j)_{j=1}^n$ -measurable as exploited before. So we have

$$V_P(\sqrt{n}\gamma_{kl,n} \mid (Y_j, Z_j)_{j=1}^n) = \frac{1}{n} \sum_{i=1}^n V_P(C_i \mid (Y_j, Z_j)_{j=1}^n) = V_P(C_i \mid (Y_j, Z_j)_{j=1}^n).$$

We compute the conditional variance to be

$$\begin{aligned} V_P(C_i \mid (Y_j, Z_j)_{j=1}^n) &= E_P \left(\varphi_k(U_{1,i})^2 \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \mid (Y_j, Z_j)_{j=1}^n \right) \\ &= \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 E_P \left(\varphi_k(U_{1,i})^2 \mid (Y_j, Z_j)_{j=1}^n \right) \\ &= \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 E_P \left(\varphi_k(U_{1,i})^2 \right) \\ &= \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \quad \text{a.s.} \end{aligned}$$

where we have used Assumption 4 (iii). We can use the the law of total variance to see that

$$\begin{aligned} V_P(\sqrt{n}\gamma_{kl,n}) &= E_P(V_P(\sqrt{n}\gamma_{kl,n} \mid (Y_j, Z_j)_{j=1}^n)) + V_P(E_P(\sqrt{n}\gamma_{kl,n} \mid (Y_j, Z_j)_{j=1}^n)) \\ &= E_P \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 + 0 = E_P \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2. \end{aligned}$$

By Lemma 30 we have that $\left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \rightarrow_P 0$ with similar arguments as before. Note that $\varphi_\ell : [0, 1] \rightarrow \mathbb{R}$ is bounded due to continuity of φ_ℓ and compactness of $[0, 1]$. Hence each term in the sequence

$$\left(\left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \right)_{i=1, \dots, n}$$

is bounded. Therefore we also have $E_P \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \rightarrow 0$. For given $\varepsilon > 0$ we have by Chebyshev's inequality that

$$P(|\sqrt{n}\gamma_{kl,n}| > \varepsilon) \leq \frac{V_P(\sqrt{n}\gamma_{kl,n})}{\varepsilon^2} = \frac{1}{\varepsilon^2} \cdot E_P \left(\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right)^2 \rightarrow 0$$

for each $P \in \mathcal{H}_0$. This shows $\sqrt{n}\gamma_n \rightarrow_P 0$. By the same argument it can be shown that $\sqrt{n}\delta_n \rightarrow_P 0$. By Slutsky's lemma we now have that

$$\sqrt{n}\hat{\rho}_n = \sqrt{n}\alpha_n + \sqrt{n}\beta_n + \sqrt{n}\gamma_n + \sqrt{n}\delta_n \Rightarrow_P \mathcal{N}(0, \Sigma \otimes \Sigma)$$

for each $P \in \mathcal{H}_0$. This shows the theorem. ■

A.9 Proof of Corollary 15

First note that Σ is a positive definite matrix as $\varphi_1, \dots, \varphi_q$ are assumed linearly independent. It thus has a positive definite matrix square root $\Sigma^{-1/2}$ satisfying $\Sigma^{-1/2}\Sigma\Sigma^{-1/2} = I$, and we have that

$$\sqrt{n}\Sigma^{-1/2}\hat{\rho}_n\Sigma^{-1/2} \Rightarrow_P \mathcal{N}(0, I \otimes I)$$

for $P \in \mathcal{H}_0$ where we have used Theorem 18. The test statistics T_n is therefore well defined and

$$nT_n = \|\sqrt{n}\Sigma^{-1/2}\hat{\rho}_n\Sigma^{-1/2}\|_F^2 \Rightarrow_P \chi_{q^2}^2$$

for $P \in \mathcal{H}_0$ by the continuous mapping theorem. \square

A.10 Proof of Corollary 17

Under Assumption 5 we have by Corollary 15 that $nT_n \Rightarrow_P \chi_{q^2}^2$. Therefore

$$\limsup_{n \rightarrow \infty} E_P(\hat{\Psi}_n) = \limsup_{n \rightarrow \infty} P(nT_n > z_{1-\alpha}) = \limsup_{n \rightarrow \infty} (1 - (F_{nT_n}(z_{1-\alpha}))) = 1 - (1 - \alpha) = \alpha.$$

because $F_{nT_n}(t) \rightarrow \Phi(t)$ as $n \rightarrow \infty$ for all $t \in \mathbb{R}$ where Φ is the distribution function of a $\chi_{q^2}^2$ -distribution and $z_{1-\alpha}$ is the $(1 - \alpha)$ -quantile of a $\chi_{q^2}^2$ -distribution. \square

A.11 Proof of Theorem 18

The proof uses the same decomposition as in the proof of Theorem 14, i.e., $\hat{\rho}_n = \alpha_n + \beta_n + \gamma_n + \delta_n$. Let us first comment on the large sample properties of α_n . Since α_n is the i.i.d. average of terms with expectation ρ for all $P \in \mathcal{P}_0$ we have that $\alpha_n \rightarrow_P \rho$ for all $P \in \mathcal{P}_0$. The term $\hat{\beta}_n$ is dealt with similarly as in the proof of Theorem 14. For fixed $k, \ell = 1, \dots, q$ we have that

$$\begin{aligned} |\beta_{k\ell, n}| &\leq \frac{1}{n} \sum_{i=1}^n \left| \varphi_k(\hat{U}_{1,i}) - \varphi_k(U_{1,i}) \right| \left| \varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i}) \right| \\ &\leq \|\varphi_k \circ \hat{F}_{X|Z}^{(n)} - \varphi_k \circ F_{X|Z}\|_\infty \cdot \|\varphi_\ell \circ \hat{F}_{Y|Z}^{(n)} - \varphi_\ell \circ F_{Y|Z}\|_\infty \in \mathcal{O}_P(g_P(n)h_P(n)) \end{aligned}$$

where we have used Lemma 30. From Assumption 5 we get that $\beta_{k\ell, n} \rightarrow_P 0$ for each $k, \ell = 1, \dots, q$, and so $\beta_n \rightarrow_P 0$ for all $P \in \mathcal{P}_0$. The terms γ_n and δ_n are analyzed similarly, so we only look at γ_n . We see that for $k, \ell = 1, \dots, q$,

$$\begin{aligned} |\gamma_{k\ell, n}| &\leq \frac{1}{n} \sum_{i=1}^n |\varphi_k(U_{1,i})| |\varphi_\ell(\hat{U}_{2,i}) - \varphi_\ell(U_{2,i})| \\ &\leq \|\varphi_k\|_\infty \cdot \|\varphi_\ell \circ \hat{F}_{Y|Z}^{(n)} - \varphi_\ell \circ F_{Y|Z}\|_\infty \in \mathcal{O}_P(h_P(n)) \end{aligned}$$

where we have used that $\|\varphi_k\|_\infty < \infty$ since $\varphi_k : [0, 1] \rightarrow \mathbb{R}$ is continuous and $[0, 1]$ is compact. Here we have used Lemma 30, and we conclude that $\gamma_{k\ell, n} \rightarrow_P 0$ due to Assumption 5, which shows that $\gamma_n \rightarrow_P 0$ for all $P \in \mathcal{P}_0$. Conclusively, we have $\hat{\rho}_n \rightarrow_P \rho$ for all $P \in \mathcal{P}_0$. \square

A.12 Proof of Corollary 19

Assume that $P \in \mathcal{A}_0$ such that $\rho_{k\ell} \neq 0$ for some $k, \ell = 1, \dots, q$. Then we have

$$T_n = \|\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2}\|_F^2 \xrightarrow{P} \|\Sigma^{-1/2} \rho \Sigma^{-1/2}\|_F^2 > 0$$

for all $P \in \mathcal{A}_0$ because $\rho \neq 0$. Here we have used Theorem 18. Therefore we obtain that

$$nT_n = n\|\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2}\|_F^2 \xrightarrow{P} \infty$$

for all $P \in \mathcal{A}_0$. This means that

$$P(nT_n > c) \rightarrow 1$$

as $n \rightarrow \infty$ for all $c \in \mathbb{R}$. From this we obtain that

$$\liminf_{n \rightarrow \infty} E_P(\hat{\Psi}_n) = \liminf_{n \rightarrow \infty} P(nT_n > z_{1-\alpha}) = 1$$

for all $\alpha \in (0, 1)$ whenever $P \in \mathcal{A}_0$. □

A.13 Proof of Proposition 20

Assume $(U_1, U_2) \perp\!\!\!\perp Z$ and $U_1 \perp\!\!\!\perp U_2$. Then it also holds that $U_1 \perp\!\!\!\perp U_2 \mid Z$, which gives $(U_1, Z) \perp\!\!\!\perp (U_2, Z) \mid Z$. More explicitly we have

$$(F_{X|Z}(X \mid Z), Z) \perp\!\!\!\perp (F_{Y|Z}(Y \mid Z), Z) \mid Z.$$

Transforming with the conditional quantile functions gives

$$Q_{X|Z}(F_{X|Z}(X \mid Z) \mid Z) \perp\!\!\!\perp Q_{Y|Z}(F_{Y|Z}(Y \mid Z) \mid Z) \mid Z.$$

Since we assume throughout the paper that the conditional distributions $X \mid Z = z$ and $Y \mid Z = z$ are continuous for each $z \in [0, 1]^d$ we get that $(X, Z) \perp\!\!\!\perp (Y, Z) \mid Z$ which reduces to $X \perp\!\!\!\perp Y \mid Z$. □

A.14 Proof of Theorem 21

We start by showing (i). Again we consider the decomposition $\hat{\rho}_n = \alpha_n + \beta_n + \gamma_n + \delta_n$ introduced in the proof of Theorem 14. By the stronger condition of Assumption 6 we immediately have that $\sqrt{n}\beta_n \rightarrow_{\mathcal{P}_0} 0$, $\sqrt{n}\gamma_n \rightarrow_{\mathcal{P}_0} 0$ and $\sqrt{n}\delta_n \rightarrow_{\mathcal{P}_0} 0$ by following the same arguments as in the proof of Theorem 14. The fact that $\sqrt{n}\alpha_n$ converges uniformly in distribution to a $\mathcal{N}(0, \Sigma \otimes \Sigma)$ -distribution over \mathcal{H}_0 follows from the fact that the distribution of $(U_{1,i}, U_{2,i})_{i=1}^n$ is unchanged whenever $P \in \mathcal{H}_0$. By Lemma 37 we have that

$$\sqrt{n}\hat{\rho}_n = \sqrt{n}\alpha_n + \sqrt{n}\beta_n + \sqrt{n}\gamma_n + \sqrt{n}\delta_n \Rightarrow_{\mathcal{H}_0} \mathcal{N}(0, \Sigma \otimes \Sigma)$$

which shows part (i) of the theorem. Next we turn to part (ii) of the theorem. Analogously to the proof of Theorem 18 we have that $\beta_n \rightarrow_{\mathcal{P}_0} 0$, $\gamma_n \rightarrow_{\mathcal{P}_0} 0$ and $\delta_n \rightarrow_{\mathcal{P}_0} 0$ under Assumption 6. Now consider writing

$$\alpha_{k\ell, n} = \frac{1}{n} \sum_{i=1}^n A_i \quad \text{where} \quad A_i = \varphi_k(U_{1,i})\varphi_\ell(U_{2,i})$$

for $k, \ell = 1, \dots, q$. Then $(A_i)_{i=1}^n$ are i.i.d. with $E_P(A_i) = \rho_{k\ell}$ and

$$V_P(A_i) = E_P(\varphi_k(U_{1,i})^2 \varphi_\ell(U_{2,i})^2) - \rho^2 \leq \|\varphi_k\|_\infty \|\varphi_\ell\|_\infty < \infty$$

for all $P \in \mathcal{P}_0$. Therefore, for given $\varepsilon > 0$, we have by Chebyshev's inequality that

$$\sup_{P \in \mathcal{P}_0} P(|\alpha_{k\ell,n} - \rho_{k\ell}| > \varepsilon) \leq \sup_{P \in \mathcal{P}_0} \frac{V_P(\frac{1}{n} \sum_{i=1}^n A_i)}{\varepsilon^2} = \sup_{P \in \mathcal{P}_0} \frac{V_P(A_i)}{n\varepsilon^2} \leq \frac{\|\varphi_1\|_\infty^2 \|\varphi_2\|_\infty^2}{n\varepsilon^2} \rightarrow 0$$

for $n \rightarrow \infty$ which shows that $\alpha_n \rightarrow_{\mathcal{P}_0} \rho$. From this we get $\hat{\rho}_n \rightarrow_{\mathcal{P}_0} \rho$ as wanted.

A.15 Proof of Corollary 22

Note that due to Theorem 21 (i) we have that $nT_n \Rightarrow_{\mathcal{H}_0} \chi_{q^2}^2$ under Assumption 6 using the same argument as in the proof of Corollary 15. Then the result is obtained by the same argument as in the proof of Corollary 17 by noting that $\sup_{P \in \mathcal{H}_0} |F_{nT_n}(t) - \Phi(t)| \rightarrow 0$ as $n \rightarrow \infty$ for all $t \in \mathbb{R}$. \square

A.16 Proof of Corollary 23

Let $\lambda > 0$ be fixed. By Theorem 21 (ii) we have $\hat{\rho}_n \rightarrow_{\mathcal{A}_\lambda} \rho$ where $|\rho_{k\ell}| > \lambda > 0$ for some $k, \ell = 1, \dots, q$. Therefore $\inf_{P \in \mathcal{A}_\lambda} |\rho_{k\ell}| \geq \lambda > 0$ and so

$$T_n = \|\Sigma^{-1/2} \hat{\rho}_n \Sigma^{-1/2}\|_F^2 \rightarrow_{\mathcal{A}_\lambda} \|\Sigma^{-1/2} \rho \Sigma^{-1/2}\|_F^2 > 0$$

since $\inf_{P \in \mathcal{A}_\lambda} |\rho_{k\ell}^P| > 0$ and $\Sigma^{-1/2}$ is positive definite. Therefore $nT_n \rightarrow_{\mathcal{A}_\lambda} \infty$, and so we have

$$\begin{aligned} \inf_{P \in \mathcal{A}_\lambda} P(nT_n > c) &= \inf_{P \in \mathcal{A}_\lambda} (1 - P(nT_n \leq c)) \\ &= - \sup_{P \in \mathcal{A}_\lambda} (P(nT_n \leq c) - 1) \rightarrow 1 \end{aligned}$$

as $n \rightarrow \infty$ for all $c \in \mathbb{R}$. From this we have

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{A}_\lambda} E_P(\hat{\Psi}_n) = \liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{A}_\lambda} P(nT_n > z_{1-\alpha}) = 1$$

for all $\alpha \in (0, 1)$. \square

Appendix B. Modes of Stochastic Convergence

Let \mathcal{M} denote some class of distributions. We start by defining the notions of small and big O in probability.

B.1 Small and big-O in probability

All sequences (a_n) and (b_n) below are assumed to be non-zero.

Definition 31 Let (X_n) and (a_n) be sequences of random variables in \mathbb{R} . If for every $\varepsilon > 0$

$$\sup_{P \in \mathcal{M}} P(|X_n/a_n| > \varepsilon) \rightarrow 0$$

for $n \rightarrow \infty$ then we say that X_n is small O of a_n in probability uniformly over \mathcal{M} and write $X_n \in o_{\mathcal{M}}(a_n)$. If for every $\varepsilon > 0$ there is $M > 0$ such that

$$\sup_{n \in \mathbb{N}} \sup_{P \in \mathcal{M}} P(|X_n/a_n| > M) < \varepsilon$$

then we say that X_n is big O of a_n in probability uniformly over \mathcal{M} and write $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$.

When $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$ we also say that X_n is stochastically bounded by a_n uniformly over \mathcal{M} . When $X_n \in o_{\mathcal{M}}(1)$ we will typically write $X_n \rightarrow_{\mathcal{M}} 0$.

Lemma 32 Let $(X_n), (a_n)$ and (b_n) be sequences of random variables in \mathbb{R} such that $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$. Then it holds that $b_n X_n \in \mathcal{O}_{\mathcal{M}}(a_n b_n)$.

Lemma 33 Assume that $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$ and $Y_n \in \mathcal{O}_{\mathcal{M}}(b_n)$. Then $X_n Y_n \in \mathcal{O}_{\mathcal{M}}(a_n b_n)$.

Lemma 34 Assume $X_n \in \mathcal{O}_{\mathcal{M}}(a_n)$ and that $a_n \in o(1)$. Then $X_n \in o_{\mathcal{M}}(1)$.

Lemma 35 Assume that $X_n \in o_{\mathcal{M}}(1)$ and that $|X_n| \leq C$ for all $n \geq 1$ for a constant C that does not depend on P . Then $\sup_{P \in \mathcal{M}} E_P |X_n| \rightarrow 0$ for $n \rightarrow \infty$.

We now turn to uniform convergence in distribution.

B.2 Uniform convergence in distribution

We follow Kasy (2019) and Bengs and Holzmann (2019).

Definition 36 Let X, X_1, X_2, \dots be real valued random variables with distribution determined by $P \in \mathcal{M}$. If it holds that

$$\sup_{P \in \mathcal{M}} |E_P(f(X_n)) - E_P(f(X))| \rightarrow 0$$

for $n \rightarrow \infty$ for all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ that are bounded and continuous, then we say that (X_n) converges uniformly in distribution to X over \mathcal{M} . In this case we write $X_n \Rightarrow_{\mathcal{M}} X$.

Lemma 37 (Uniform Slutsky's Lemma) Assume that $X_n \Rightarrow_{\mathcal{M}} X$ and that $Y_n \rightarrow_{\mathcal{M}} 0$. Then $X_n + Y_n \Rightarrow_{\mathcal{M}} X$.

Proof See Bengs and Holzmann (2019) Theorem 6.3. ■

References

- Alexandre Belloni and Victor Chernozhukov. ℓ_1 -penalized quantile regression in high-dimensional sparse models. *The Annals of Statistics*, 39(1):82–130, 2011.
- Alexandre Belloni, Victor Chernozhukov, and Kengo Kato. Valid post-selection inference in high-dimensional approximately sparse quantile regression models. *Journal of the American Statistical Association*, 114(526):749–758, 2019.
- Viktor Bengs and Hajo Holzmann. Uniform approximation in classical weak convergence theory. *arXiv preprint arXiv:1903.09864*, 2019.
- Wicher Bergsma. *Testing conditional independence for continuous random variables*. Eu-random, 2004.
- Wicher Bergsma. Nonparametric testing of conditional independence by means of the partial copula. *SSRN Electronic Journal*, 01 2011.
- Taoufik Bouezmarni, Jeroen VK Rombouts, and Abderrahim Taamouti. Nonparametric copula-based test for conditional independence with applications to Granger causality. *Journal of Business & Economic Statistics*, 30(2):275–287, 2012.
- Victor Chernozhukov, Iván Fernández-Val, and Alfred Galichon. Quantile and probability curves without crossing. *Econometrica*, 78(3):1093–1125, 2010.
- A Philip Dawid. Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(1):1–15, 1979.
- Uwe Einmahl and David M. Mason. Uniform in bandwidth consistency of kernel-type function estimators. *The Annals of Statistics*, 33(3):1380 – 1403, 2005.
- Jianqing Fan, Yang Feng, and Lucy Xia. A projection-based conditional dependence measure with applications to high-dimensional undirected graphical models. *Journal of Econometrics*, 218(1):119–139, 2020.
- Irène Gijbels, Marek Omelka, and Noël Veraverbeke. Estimation of a copula when a covariate affects only marginal distributions. *Scandinavian Journal of Statistics*, 42(4): 1109–1126, 2015.
- Ingrid Hobæk Haff and Johan Segers. Nonparametric estimation of pair-copula constructions with the empirical pair-copula. *Computational Statistics & Data Analysis*, 84:1–13, 2015.
- Naftali Harris and Mathias Drton. PC algorithm for nonparanormal graphical models. *Journal of Machine Learning Research*, 14(1):3365–3383, 2013.
- Xuming He. Quantile curves without crossing. *The American Statistician*, 51(2):186–192, 1997.
- Maximilian Kasy. Uniformity and the delta method. *Journal of Econometric Methods*, 8(1), 2019.

- Roger Koenker. *Quantile Regression*. Econometric Society Monographs. Cambridge University Press, 2005.
- Roger Koenker. *quantreg: Quantile Regression*, 2021. URL <https://CRAN.R-project.org/package=quantreg>. R package version 5.83.
- Roger Koenker, Victor Chernozhukov, Xuming He, and Limin Peng. *Handbook of Quantile Regression*. CRC press, 2017.
- Steffen L. Lauritzen. *Graphical Models*, volume 17 of *Oxford Statistical Science Series*. The Clarendon Press, Oxford University Press, New York, 1996. Oxford Science Publications.
- Qi Liu, Chun Li, Valentine Wanga, and Bryan E Shepherd. Covariate-adjusted Spearman’s rank correlation with probability-scale residuals. *Biometrics*, 74(2):595–605, 2018.
- Rohit K Patra, Bodhisattva Sen, and Gábor J Székely. On a nonparametric notion of residual and its applications. *Statistics & Probability Letters*, 109:208–213, 2016.
- Andrew J Patton. Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2):527–556, 2006.
- Judea Pearl. *Causality*. Cambridge University Press, second edition, 2009.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL <https://www.R-project.org/>.
- Joseph D Ramsey. A scalable conditional independence test for nonlinear, non-Gaussian data. *arXiv preprint arXiv:1401.5031*, 2014.
- Murray Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.
- Rajen D. Shah and Jonas Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3):1514 – 1538, 2020.
- Kyungchul Song. Testing conditional independence via Rosenblatt transforms. *The Annals of Statistics*, 37(6B):4011–4045, 2009.
- Fabian Spanhel and Malte S Kurz. The partial vine copula: A dependence measure and approximation based on the simplifying assumption. *arXiv preprint arXiv:1510.06971*, 2015.
- Fabian Spanhel and Malte S. Kurz. The partial copula: Properties and associated dependence measures. *Statistics & Probability Letters*, 119:76–83, 2016.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, second edition, 2000.
- Winfried Stute et al. On almost sure convergence of conditional empirical distribution functions. *The Annals of Probability*, 14(3):891–901, 1986.