

THE HARDNESS OF CONDITIONAL INDEPENDENCE TESTING AND THE GENERALISED COVARIANCE MEASURE

BY RAJEN D. SHAH¹ AND JONAS PETERS²

¹*Statistical Laboratory, University of Cambridge, r.shah@statslab.cam.ac.uk*

²*Department of Mathematical Sciences, University of Copenhagen, jonas.peters@math.ku.dk*

It is a common saying that testing for conditional independence, that is, testing whether two random vectors X and Y are independent, given Z , is a hard statistical problem if Z is a continuous random variable (or vector). In this paper, we prove that conditional independence is indeed a particularly difficult hypothesis to test for. Valid statistical tests are required to have a size that is smaller than a pre-defined significance level, and different tests usually have power against a different class of alternatives. We prove that a valid test for conditional independence does not have power against *any* alternative.

Given the nonexistence of a uniformly valid conditional independence test, we argue that tests must be designed so their suitability for a particular problem may be judged easily. To address this need, we propose in the case where X and Y are univariate to nonlinearly regress X on Z , and Y on Z and then compute a test statistic based on the sample covariance between the residuals, which we call the generalised covariance measure (GCM). We prove that validity of this form of test relies almost entirely on the weak requirement that the regression procedures are able to estimate the conditional means X given Z , and Y given Z , at a slow rate. We extend the methodology to handle settings where X and Y may be multivariate or even high dimensional. While our general procedure can be tailored to the setting at hand by combining it with any regression technique, we develop the theoretical guarantees for kernel ridge regression. A simulation study shows that the test based on GCM is competitive with state of the art conditional independence tests. Code is available as the R package `GeneralisedCovarianceMeasure` on CRAN.

1. Introduction. Conditional independences lie at the heart of several fundamental concepts such as sufficiency [21] and ancillarity [22, 23]; see also Jensen and Sørensen [28]. Dawid [17] states that “many results and theorems concerning these concepts are just applications of some simple general properties of conditional independence”. During the last few decades, conditional independence relations have played an increasingly important role in computational statistics, too, since they are the building blocks of graphical models [30, 32, 37].

Estimating conditional independence graphs has been of great interest in high-dimensional statistics, particularly in biomedical applications (e.g., [18, 35]). To estimate the conditional independence graph corresponding to a random vector $X \in \mathbb{R}^d$, edges may be placed between vertices corresponding to X_j and X_k if a test for whether X_j is conditionally independent of X_k given all other variables indicates rejection.

Conditional independence tests also play a key role in causal inference. Constraint-based or independence-based methods [37, 39, 48] apply a series of conditional independence tests

Received April 2018; revised April 2019.

MSC2010 subject classifications. Primary 62G10; secondary 62G08.

Key words and phrases. Conditional independence, hypothesis testing, testability, wild bootstrap, kernel ridge regression.

in order to learn causal structure from observational data. The recently introduced invariant prediction methodology [26, 38] aims to estimate for a specified target variable Y , the set of causal variables among potential covariates X_1, \dots, X_d . Given data from different environments labelled by a variable E , the method involves testing for each subset $S \subseteq \{1, \dots, d\}$, the null hypothesis that the environment E is conditionally independent of Y , given X_S .

Given the importance of conditional independence tests in modern statistics, there has been great deal of work devoted to developing conditional independence tests; we review some important examples of tests in Section 1.3. However, one issue that conditional independence tests appear to suffer from is that they can fail to control the type I error in finite samples, which can have important consequences in downstream analyses.

In the first part of this paper, we prove this failure of type I error control is in fact unavoidable: conditional independence is not a testable hypotheses. To fix ideas, consider n i.i.d. observations corresponding to a triple of random variables (X, Y, Z) where it is desired to test whether X is conditional independent of Y given Z . We show that provided the joint distribution of $(X, Y, Z) \in \mathbb{R}^{d_X+d_Y+d_Z}$ is absolutely continuous with respect to Lebesgue measure, any test based on the data whose size is less than a pre-specified level α , has no power; more precisely, there is no alternative, for which the test has power more than α . This result is perhaps surprising as it is in stark contrast to unconditional independence testing, for which permutation tests allow for the correct calibration of any hypothesis test. Our result implies that in order to perform conditional independence testing, some domain knowledge is required in order to select an appropriate conditional independence test for the particular problem at hand. This would appear to be challenging in practice, as the validity of conditional independence tests typically rests on properties of the entire joint distribution of the data, which may be hard to model.

Our second main contribution aims to alleviate this issue by providing a family of conditional independence tests whose validity relies on estimating the conditional expectations $\mathbb{E}(X|Z = z)$ and $\mathbb{E}(Y|Z = z)$ via regressions, in the setting where $d_X = d_Y = 1$. These need to be estimated sufficiently well using the data such that the product of mean squared prediction errors from the two regressions is $o(n^{-1})$. This is a relatively mild requirement that allows for settings where the conditional expectations are as general as Lipschitz functions, for example, and also encompasses settings where Z is high dimensional but the conditional expectations have more structure.

Our test statistic, which we call the *generalised covariance measure* (GCM) is based on a suitably normalised version of the empirical covariance between the residual vectors from the regressions. The practitioner is free to choose the regression methods that appear most suitable for the problem of interest. Although domain knowledge is still required to make an appropriate choice, selection of regression methods is a problem statisticians are more familiar with. We also extend the GCM to handle settings where X and Y are potentially high dimensional, though in this case our proof of the validity of the test additionally requires the errors $X_j - \mathbb{E}(X_j|Z)$ and $Y_k - \mathbb{E}(Y_k|Z)$ to obey certain moment restrictions for $j = 1, \dots, d_X$ and $k = 1, \dots, d_Y$ and slightly faster rates of convergence for the prediction errors.

As an example application of our results on the GCM, we consider the case where the regressions are performed using kernel ridge regression, and show that provided the conditional expectations are contained in a reproducing kernel Hilbert space, our test statistic has a tractable limit distribution.

The rest of the paper is organised as follows. In Sections 1.1 and 1.2, we first formalise the notion of conditional independence and relevant concepts related to statistical hypothesis testing. In Section 1.3, we review some popular conditional independence tests, after which we set out some notation used throughout the paper. In Section 2, we present our main result on the hardness of conditional independence testing. We introduce the generalised covariance measure in Section 3 first treating the univariate case with $d_X = d_Y = 1$ before

extending ideas to the potentially high-dimensional case. In Section 4, we apply the theory and methodology of the previous section to study that particular example of generalised covariance measures based on kernel ridge regression. We present numerical experiments in Section 5 and conclude with a discussion in Section 6. All proofs are deferred to the Appendix and Supplementary Material [47].

1.1. *Conditional independence.* Let us consider three random vectors X, Y and Z taking values in $\mathbb{R}^{d_x}, \mathbb{R}^{d_y}$ and \mathbb{R}^{d_z} , respectively, and let us assume, for now that their joint distribution is absolutely continuous with respect to Lebesgue measure with density p . For our deliberations, only the continuity in Z is necessary; see Remark 4. We say that X is conditionally independent of Y given Z and write

$$X \perp\!\!\!\perp Y \mid Z$$

if for all x, y, z with $p(z) > 0$, we have $p(x, y|z) = p(x|z)p(y|z)$; see, for example, Dawid [17]. Here and below, statements involving densities should be understood to hold (Lebesgue) almost everywhere. We now discuss an equivalent formulation of conditional independence that has given rise to several hypothesis tests, including the generalised covariance measure proposed in this paper. Let therefore $L^2_{X,Z}$ denote the space of all functions $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}$ such that $\mathbb{E}f(X, Z)^2 < \infty$ and define $L^2_{Y,Z}$ analogously. Daudin [16] proves that X and Y are conditionally independent given Z if and only if

$$(1) \quad \mathbb{E}f(X, Z)g(Y, Z) = 0$$

for all functions $f \in L^2_{X,Z}$ and $g \in L^2_{Y,Z}$ such that $\mathbb{E}[f(X, Z)|Z] = 0$ and $\mathbb{E}[g(Y, Z)|Z] = 0$, respectively.

This may be viewed as an extension of the fact that for one-dimensional X and Y , the partial correlation coefficient $\rho_{X,Y|Z}$ (the correlation between residuals of linear regressions of X on Z and Y on Z) is 0 if and only if $X \perp\!\!\!\perp Y \mid Z$ in the case where (X, Y, Z) are jointly Gaussian.

1.2. *Statistical hypothesis testing and notation.* We now introduce some notation and relevant concepts related to statistical hypothesis testing. In order to deal with composite null hypotheses where the probability of rejection must be controlled under a variety of different distributions for the data to which our test is applied, we introduce the following notation. We will write $\mathbb{E}_P(\cdot)$ for expectations of random variables whose distribution is determined by P , and similarly $\mathbb{P}_P(\cdot) = \mathbb{E}_P \mathbb{1}_{\{\cdot\}}$.

Let \mathcal{P} be a potentially composite null hypothesis consisting of a collection of distributions for (X, Y, Z) . For $i = 1, 2, \dots$ let $(x_i, y_i, z_i) \in \mathbb{R}^{d_x+d_y+d_z}$ be i.i.d. copies of (X, Y, Z) and let $\mathbf{X}^{(n)} \in \mathbb{R}^{d_x \cdot n}$, $\mathbf{Y}^{(n)} \in \mathbb{R}^{d_y \cdot n}$ and $\mathbf{Z}^{(n)} \in \mathbb{R}^{d_z \cdot n}$ be matrices with i th rows x_i, y_i and z_i , respectively. Let ψ_n be a potentially randomised test that can be applied to the data $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$; formally,

$$\psi_n : \mathbb{R}^{(d_x+d_y+d_z) \cdot n} \times [0, 1] \rightarrow \{0, 1\}$$

is a measurable function whose last argument is reserved for a random variable $U \sim U[0, 1]$ independent of the data which is responsible for the randomness of the test.

Given a sequence of tests $(\psi_n)_{n=1}^\infty$, the following validity properties will be of interest; note the particular names given to these properties differ in literature. Given a level $\alpha \in (0, 1)$ and null hypothesis \mathcal{P} , we say that the test ψ_n has

$$\text{valid level at sample size } n \text{ if } \sup_{P \in \mathcal{P}} \mathbb{P}_P(\psi_n = 1) \leq \alpha,$$

where the left-hand side is the *size* of the test; the sequence $(\psi_n)_{n=1}^\infty$ has

$$\text{uniformly asymptotic level if } \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \mathbb{P}_P(\psi_n = 1) \leq \alpha,$$

$$\text{pointwise asymptotic level if } \sup_{P \in \mathcal{P}} \limsup_{n \rightarrow \infty} \mathbb{P}_P(\psi_n = 1) \leq \alpha.$$

In practice, we would like a test to have at least uniformly asymptotic level. Otherwise, even for an arbitrarily large sample size n , there can exist null distributions for which the size exceeds the nominal level by some fixed amount.

Given a sequence of tests $(\psi_n)_{n=1}^\infty$ each with valid level $\alpha \in (0, 1)$ and alternative hypotheses \mathcal{Q} , it is desirable for the power to be large uniformly over \mathcal{Q} , and to have $\inf_{Q \in \mathcal{Q}} \mathbb{P}_Q(\psi_n = 1) \rightarrow 1$. In standard parametric settings, we can certainly achieve this for any fixed alternative hypothesis and indeed have uniform power against a sequence of \sqrt{n} alternatives. Nonparametric problems are much harder and when \mathcal{Q} contains all distributions outside a small fixed total variation (TV) neighbourhood of \mathcal{P} , we have

$$\liminf_{n \rightarrow \infty} \inf_{Q \in \mathcal{Q}} \mathbb{P}_Q(\psi_n = 1) < 1,$$

[33], [4], Proposition 2, Theorem 3. To achieve power tending to 1, we need to restrict \mathcal{Q} by imposing certain smoothness conditions for example [3].

A class of even harder hypothesis testing problems may be defined as those where no test with valid level achieves power at any alternative, so $\sup_{Q \in \mathcal{Q}} \mathbb{P}_Q(\psi_n = 1) \leq \alpha$. In other words, for all n , tests ψ_n and alternative distributions $Q \in \mathcal{Q}$, we have

$$\mathbb{P}_Q(\psi_n = 1) \leq \sup_{P \in \mathcal{P}} \mathbb{P}_P(\psi_n = 1).$$

The hypothesis testing problem defined by the pair $(\mathcal{P}, \mathcal{Q})$ is then said to be *untestable* [19]. In order to have power at even a single alternative, we need to restrict the null \mathcal{P} in some way. One of the main results of this paper is that conditional independence is untestable.

1.3. *Related work.* Our hardness result for conditional independence contributes to an important literature on impossibility results in statistics and econometrics starting with the work of Bahadur and Savage [2], which shows that there is no nontrivial test for whether a distribution has mean zero. Canay et al. [11] shows that certain problems arising in the context of identification of some nonparametric models are not testable. In these examples, the null hypothesis is dense with respect to the TV metric in the alternative hypothesis, a property which implies untestability [45]. Interestingly, our Proposition 5 shows that conditional independence testing is qualitatively different in that some distributions in the alternative are in fact well-separated from the null. It has been suggested for some time that conditional independence testing is a hard problem (see, e.g., [5], and several talks given by Bernhard Schölkopf). To the best of our knowledge, the conjecture that conditional independence is not testable (cf. Corollary 3 with $M = \infty$) is due to Arthur Gretton. We also note that when the conditional distribution of X given Z is known, conditional independence is testable (e.g., [7]).

We now briefly review several tests for conditional independence that bear some relation to our proposal here.

Extensions of partial correlation. Ramsey [40] suggests regressing X on Z and Y on Z and then testing for independence between the residuals. Fan et al. [20] consider this approach in the setting where Z is potentially high-dimensional and under the null hypothesis of $X \perp\!\!\!\perp Y \mid Z$, $X = Z^T \beta_X + \varepsilon_X$, $Y = Z^T \beta_Y + \varepsilon_Y$ with $\varepsilon_X \perp\!\!\!\perp Z$ and $\varepsilon_Y \perp\!\!\!\perp Z$. The following simple example however indicates where such methods can fail.

EXAMPLE 1. Define N_X, N_Y and Z to be i.i.d. random variables with distribution $\mathcal{N}(0, 1)$ and define $X := Z \cdot N_X, Y := Z \cdot N_Y$. This implies $X \perp\!\!\!\perp Y \mid Z$. Since $\mathbb{E}[X|Z] = \mathbb{E}[Y|Z] = 0$, the (population) residuals equal $R_1 := Z \cdot N_X$ and $R_2 := Z \cdot N_Y$; they are uncorrelated but not independent since, for example, $\text{cov}(R_1^2, R_2^2) \neq 0$. Consider regressing X on Z and Y on Z , and then testing for independence of the residuals. If the regression method outputs the true conditional means and the independence test has power against the alternative $\text{cov}(R_1^2, R_2^2) \neq 0$, the method will falsely reject the null hypothesis of conditional independence with large probability.

Kernel-based conditional independence tests. The Hilbert–Schmidt independence criterion (HSIC) equals the square of the Hilbert–Schmidt norm of the cross-covariance operator, and is used in unconditional independence testing [25]. Fukumizu et al. [24] extend this idea to conditional independence testing. To construct a test for continuous variables Z , their work requires clustering of the values of Z and permuting X and Y values within the same cluster component. Another extension is proposed by Zhang et al. [51]. Their kernel conditional independence (KCI) test is stated to yield pointwise asymptotic level control.

Estimation of expected conditional covariance. Though typically not thought of as conditional independence tests, there are several approaches to estimating the expected conditional covariance functional $\mathbb{E}\text{cov}(X, Y|Z)$ in the semiparametric statistics literature [14, 36, 42]. From (1), we see these may be used as conditional independence tests and indeed the GCM test we propose falls under this category. We delay further discussion of such methods to Section 3.1.2.

1.4. *Notation.* We now introduce some notation used throughout the paper. If $(V_{P,n})_{n \in \mathbb{N}, P \in \mathcal{P}}$ is a family of sequences of random variables whose distributions are determined by $P \in \mathcal{P}$, we use $V_{P,n} = o_{\mathcal{P}}(1)$ and $V_{P,n} = O_{\mathcal{P}}(1)$ to mean respectively that for all $\epsilon > 0$,

$$\sup_{P \in \mathcal{P}} \mathbb{P}_P(|V_{P,n}| > \epsilon) \rightarrow 0, \quad \text{and}$$

$$\text{there exists } M > 0 \text{ such that } \sup_{n \in \mathbb{N}} \sup_{P \in \mathcal{P}} \mathbb{P}_P(|V_{P,n}| > M) < \epsilon.$$

If $(W_{P,n})_{n \in \mathbb{N}, P \in \mathcal{P}}$ is a further family of sequences of random variables, $V_{P,n} = o_{\mathcal{P}}(W_{P,n})$ and $V_{P,n} = O_{\mathcal{P}}(W_{P,n})$ mean $V_{P,n} = W_{P,n}R_{P,n}$ and respectively that $R_{P,n} = o_{\mathcal{P}}(1)$ and $R_{P,n} = O_{\mathcal{P}}(1)$. If \mathbf{A} is a $c \times d$ matrix, then \mathbf{A}_j denotes the j th column of \mathbf{A} , $j \in \{1, \dots, d\}$.

2. No-free-lunch in conditional independence testing. In this section, we show that, under certain conditions, no nontrivial test for conditional independence with valid level exists. To state our result, we introduce the following subsets of \mathcal{E}_0 defined to be the set of all distributions for (X, Y, Z) absolutely continuous with respect to Lebesgue measure.

Let $\mathcal{P}_0 \subset \mathcal{E}_0$ be the subset of distributions under which $X \perp\!\!\!\perp Y \mid Z$. Further, for any $M \in (0, \infty]$, let $\mathcal{E}_{0,M} \subseteq \mathcal{E}_0$ be the subset of all distributions with support contained strictly within a ℓ_∞ ball of radius M . Here, we take $\mathcal{E}_{0,\infty} = \mathcal{E}_0$. We also define $\mathcal{Q}_0 = \mathcal{E}_0 \setminus \mathcal{P}_0$ and set $\mathcal{P}_{0,M} = \mathcal{E}_{0,M} \cap \mathcal{P}_0$, and $\mathcal{Q}_{0,M} = \mathcal{E}_{0,M} \cap \mathcal{Q}_0$. Consider the setup of Section 1.2 with null hypothesis $\mathcal{P} = \mathcal{P}_{0,M}$. Our first result shows that with this null hypothesis, any test ψ_n with valid level at sample size n has no power against any alternative.

THEOREM 2 (No-free-lunch). *Given any $n \in \mathbb{N}, \alpha \in (0, 1), M \in (0, \infty]$, and any potentially randomised test ψ_n that has valid level α for the null hypothesis $\mathcal{P}_{0,M}$, we have that $\mathbb{P}_{\mathcal{Q}}(\psi_n = 1) \leq \alpha$ for all $\mathcal{Q} \in \mathcal{Q}_{0,M}$. Thus ψ_n cannot have power against any alternative.*

A proof is given in the [Appendix](#). Note that taking M to be finite ensures all the random vectors (x_i, y_i, z_i) are bounded. Thus, for example, averages will converge in distribution to Gaussian limits uniformly over $\mathcal{P}_{0,M}$; however, as the result shows, this does not help in the construction of a nontrivial test for conditional independence. An immediate corollary to Theorem 2 is that there is no nontrivial test for conditional independence with uniformly asymptotic level.

COROLLARY 3. *For all $M \in (0, \infty]$ and for any sequence $(\psi_n)_{n=1}^\infty$ of tests, we have*

$$\sup_{Q \in \mathcal{Q}_{0,M}} \limsup_{n \rightarrow \infty} \mathbb{P}_Q(\psi_n = 1) \leq \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_{0,M}} \mathbb{P}_P(\psi_n = 1).$$

This result is in stark contrast to unconditional independence testing, where a permutation test can always be used to control the size of any testing procedure. As a consequence, there exist tests with valid level at sample size n and nontrivial power. For example, Hoeffding [27] introduces a rank-based test in the case of univariate random variables and proves that it maintains uniformly asymptotic level and has asymptotic power against each fixed alternative. For the multivariate case, Berrett and Samworth [6] consider a test based on mutual information and prove level guarantees, as well as uniform power results against a wide class of alternatives. Thus while independence testing remains a hard problem in that it is only possible to have uniform power against certain subsets of alternatives, this is different to conditional independence testing where we can only hope to control the *size* uniformly over certain subsets of the *null* hypothesis.

REMARK 4. Inspection of the proof shows that Theorem 2 also holds in the case where the variables X and Y have marginal distributions that are absolutely continuous with respect to counting measure, for example. Theorem 2 therefore contains an impossibility result for testing the equality of two conditional distributions (by taking Y to be an indicator specifying the distribution). The continuity of Z , however, is necessary. If Z only takes values in $\{1, 2\}$, for example, one can reduce the problem of conditional independence testing to unconditional independence testing by combining the tests for $X \perp\!\!\!\perp Y \mid Z = 1$ and $X \perp\!\!\!\perp Y \mid Z = 2$.

The null hypothesis being dense with respect to TV distance among the alternative hypothesis is a sufficient condition for the problem to be untestable [45]. Proposition 5, proved in the supplementary material, illustrates that this is not the case here: at least for $M \in (0, \infty)$, there exists an alternative, for which there is no distribution from the null that is arbitrarily close.

PROPOSITION 5. *For $P, Q \in \mathcal{E}_0$, the total variation distance is given by*

$$\|P - Q\|_{\text{TV}} := \sup_{A \in \mathcal{B}} |\mathbb{P}_P((X, Y, Z) \in A) - \mathbb{P}_Q((X, Y, Z) \in A)|,$$

where \mathcal{B} is the Borel σ -algebra on $\mathbb{R}^{d_X+d_Y+d_Z}$. For each $M \in (0, \infty)$, there exists $Q \in \mathcal{Q}_{0,M}$ satisfying

$$\inf_{P \in \mathcal{P}_{0,M}} \|P - Q\|_{\text{TV}} \geq 1/24.$$

In Proposition 16 in the [Appendix](#), we also show that the null and alternative hypotheses are well separated in the sense of KL divergence. On the other hand, it is known that if a problem is untestable, the convex closure of the null must contain the alternative [31], [8], Theorem 5 and Corollary 1, resp. The problem of conditional independence testing therefore

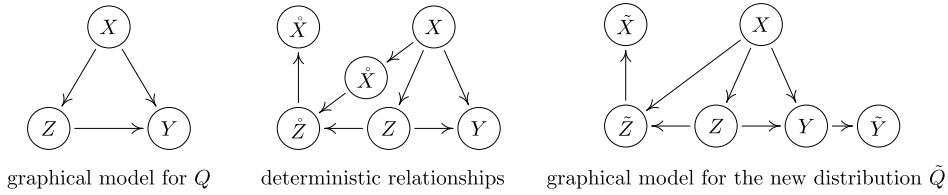


FIG. 1. Illustration of the main idea of the proof of Theorem 2. Left: We start with a distribution $Q \in \mathcal{Q}_{0,M}$ over (X, Y, Z) . In general, the Q is Markov only to a fully connected graphical model. Middle: After discretising X to \hat{X} , we are able to “hide” variable \hat{X} in $\tilde{Z} = f(Z, \hat{X})$ such that variable \tilde{Z} is close to Z in ℓ_∞ -norm and \hat{X} can be reconstructed from \tilde{Z} . Thus, \hat{X} , does not contain any “additional information” about Y when conditioning on \tilde{Z} . Right: We then consider noisy versions of the variables to guarantee that the new distribution P (over $\hat{X}, \tilde{Y}, \tilde{Z}$) is absolutely continuous with respect to Lebesgue measure, and has $\hat{X} \perp\!\!\!\perp \tilde{Y} \mid \tilde{Z}$. (The noise in \tilde{Z} is such that it still allows us to reconstruct \hat{X} from \tilde{Z} .)

has the interesting property of the null being separated from the alternative, but its convex hull is TV-dense in the alternative.

A practical implication of the negative result of Theorem 2 is that domain knowledge is needed to select a conditional independence test appropriate for the data at hand. However, guessing the form of the entire joint distribution in order to apply a test with the appropriate type I error control seems challenging. In Section 3, we introduce a form of test that instead relies on selecting regression methods that have sufficiently low prediction error when regressing $\mathbf{Y}^{(n)}$ and $\mathbf{X}^{(n)}$ on $\mathbf{Z}^{(n)}$, thereby converting the problem of finding an appropriate test to the more familiar task of prediction. Before discussing this methodology, we first sketch some of the main ideas of the proof of Theorem 2 below.

2.1. *Proof ideas of Theorem 2.* Consider the case where $d_X = d_Y = d_Z = 1$ and where the test is required to be nonrandomised. First, suppose that for $Q \in \mathcal{Q}_{0,M}$, we have a test with rejection region $R := \psi_n^{-1}(1) \subseteq \mathbb{R}^{3 \cdot n}$ such that $\mathbb{P}_Q((\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)}) \in R) > \alpha$. Let us suppose for now that R has the particularly simple form of a finite union of boxes. Our argument now proceeds by showing that one can construct a distribution $P \in \mathcal{P}_{0,M}$ from the null such that there is a coupling of P^n and Q^n where samples from each distribution are ϵ close in ℓ_∞ -norm. For a sufficiently small ϵ , we will have $\mathbb{P}_P((\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)}) \in R) > \alpha$ as well, giving the result.

Figure 1 sketches the main components in our construction of P , which is laid out formally in Lemmas 13 and 14 in the Appendix. The key idea is as follows. Given $(X, Y, Z) \sim P$, we consider a binary expansion of (X, Y, Z) , which we truncate at some point to obtain $(\hat{X}, \hat{Y}, \hat{Z})$. We then concatenate the digits of \hat{X} and \hat{Z} placing the former at the end of the binary expansion, thereby embedding \hat{X} within \hat{Z} . This way, \hat{X} can be reconstructed from \hat{Z} , and adding noise gives a distribution that is absolutely continuous with respect to Lebesgue measure. By making the truncation point sufficiently far down the expansions, we can ensure the ϵ proximity required.

For a general rejection region, we first approximate it using a finite union of boxes R^\sharp . The argument sketched above gives us $\mathbb{P}_P((\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)}) \in R^\sharp) > \alpha$, but in order to conclude the final result, we must argue that we can construct P such that $\mathbb{P}_P((\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)}) \in R^\sharp \setminus R)$ is sufficiently small. To do this, we consider a large number of potential embeddings for which the supports of the resulting distributions have little overlap. Using a probabilistic argument, we can then show that at least one embedding yields a distribution P such that the above is satisfied.

3. The generalised covariance measure. We have seen how conditional independence testing is not possible without restricting the null hypothesis. In this section, we give a general

construction for a conditional independence test based on regression procedures for regressing $\mathbf{Y}^{(n)}$ and $\mathbf{X}^{(n)}$ on $\mathbf{Z}^{(n)}$. In the case where $d_X = d_Y = 1$, which we treat in the next section, the basic form of our test statistic is a normalised covariance between the residuals from these regressions. Because of this, we call our test statistic the *generalised covariance measure* (GCM). In Section 3.2, we show how to extend the approach to handle cases where more generally $d_X, d_Y \geq 1$.

3.1. *Univariate X and Y.* Given a distribution P for (X, Y, Z) , we can always decompose

$$X = f_P(Z) + \varepsilon_P, \quad Y = g_P(Z) + \xi_P,$$

where $f_P(z) = \mathbb{E}_P(X|Z = z)$ and $g_P(z) = \mathbb{E}_P(Y|Z = z)$. Similarly, for $i = 1, 2, \dots$, we define $\varepsilon_{P,i}$ and $\xi_{P,i}$ by $x_i - f_P(z_i)$ and $y_i - g_P(z_i)$, respectively. Also let $u_P(z) = \mathbb{E}_P(\varepsilon_P^2|Z = z)$ and $v_P(z) = \mathbb{E}_P(\xi_P^2|Z = z)$.

Let $\hat{f}^{(n)}$ and $\hat{g}^{(n)}$ be estimates of the conditional expectations f_P and g_P formed, for example, by regressing $\mathbf{X}^{(n)}$ and $\mathbf{Y}^{(n)}$ on $\mathbf{Z}^{(n)}$. For $i = 1, \dots, n$, we compute the product between residuals from the regressions:

$$(2) \quad R_i = \{x_i - \hat{f}(z_i)\}\{y_i - \hat{g}(z_i)\}.$$

Here, and in what follows, we have sometimes suppressed dependence on n and P for simplicity of presentation. We then define $T^{(n)}$ to be a normalised sum of the R_i 's:

$$(3) \quad T^{(n)} = \frac{\sqrt{n} \cdot \frac{1}{n} \sum_{i=1}^n R_i}{\left(\frac{1}{n} \sum_{i=1}^n R_i^2 - \left(\frac{1}{n} \sum_{r=1}^n R_r\right)^2\right)^{1/2}} =: \frac{\tau_N^{(n)}}{\tau_D^{(n)}}.$$

Our final test can be based on $|T^{(n)}|$ with large values suggesting rejection. Note that the introduction of notation for the numerator and denominator in the definition of $T^{(n)}$ are for later use in Theorem 8.

In the case where \hat{f} and \hat{g} are formed through linear regressions, the test is similar to one based on partial correlation, and would be identical were the denominator in (3) to be replaced by the the product of the empirical standard deviations of the vectors $(x_i - \hat{f}(z_i))_{i=1}^n$ and $(y_i - \hat{g}(z_i))_{i=1}^n$. This approach however would fail for Example 1 despite f and g being linear (in fact both equal to the zero function) as the product of the variances of the residuals would not in general equal the variance of their product. Indeed, the reader may convince herself using `pcor.test` from the R package `ppcor` [29], for example, that common tests for vanishing partial correlation do not yield the correct size in this case.

The following result gives conditions under which when the null hypothesis of conditional independence holds, we can expect the asymptotic distribution of $T^{(n)}$ to be a standard normal.

THEOREM 6. *Define the following quantities:*

$$A_f := \frac{1}{n} \sum_{i=1}^n \{f_P(z_i) - \hat{f}(z_i)\}^2, \quad B_f := \frac{1}{n} \sum_{i=1}^n \{f_P(z_i) - \hat{f}(z_i)\}^2 v_P(z_i),$$

$$A_g := \frac{1}{n} \sum_{i=1}^n \{g_P(z_i) - \hat{g}(z_i)\}^2, \quad B_g := \frac{1}{n} \sum_{i=1}^n \{g_P(z_i) - \hat{g}(z_i)\}^2 u_P(z_i).$$

We have the following results:

(i) If for $P \in \mathcal{P}_0$, $A_f A_g = o_P(n^{-1})$, $B_f = o_P(1)$, $B_g = o_P(1)$ and also $0 < \mathbb{E}_P(\varepsilon_P^2 \xi_P^2) < \infty$, then

$$\sup_{t \in \mathbb{R}} |\mathbb{P}_P(T^{(n)} \leq t) - \Phi(t)| \rightarrow 0.$$

(ii) Let $\mathcal{P} \subset \mathcal{P}_0$ be a class of distributions such that $A_f A_g = o_P(n^{-1})$, $B_f = o_P(1)$, $B_g = o_P(1)$. If in addition $\inf_{P \in \mathcal{P}} \mathbb{E}(\varepsilon_P^2 \xi_P^2) \geq c_1$ and $\sup_{P \in \mathcal{P}} \mathbb{E}_P\{|\varepsilon_P \xi_P|^{2+\eta}\} \leq c_2$ for some $c_1, c_2 > 0$ and $\eta > 0$, then

$$\sup_{P \in \mathcal{P}} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(T^{(n)} \leq t) - \Phi(t)| \rightarrow 0.$$

REMARK 7. Applying the Cauchy–Schwarz inequality and Markov’s inequality, we see the requirement that $A_f A_g = o_P(n^{-1})$ is fulfilled if

$$(4) \quad n \mathbb{E}_P \left(\frac{1}{n} \sum_{i=1}^n \{f_P(z_i) - \hat{f}(z_i)\}^2 \right) \mathbb{E}_P \left(\frac{1}{n} \sum_{i=1}^n \{g_P(z_i) - \hat{g}(z_i)\}^2 \right) \rightarrow 0.$$

Thus if in addition we have $\mathbb{E}_P B_f, \mathbb{E}_P B_g \rightarrow 0$, this is sufficient for all conditions required in (i) to hold.

If $\mathbb{E}_P B_f, \mathbb{E}_P B_g$ and the left-hand side of (4) converges to 0 uniformly over all $P \in \mathcal{P}$, then the conditions in (ii) will hold provided the moment condition on $\varepsilon_P \xi_P$ is also satisfied.

A proof is given in the supplementary material. We see that under conditions largely to do with the mean squared prediction error (MSPE) of \hat{f} and \hat{g} , $T^{(n)}$ can be shown to be asymptotically standard normal (i), and if the prediction error is uniformly small, the convergence to the Gaussian limit is correspondingly uniform (ii). A key point is that the requirement on the predictive properties of \hat{f} and \hat{g} is reasonably weak: for example, provided their MSPEs are $o(n^{-1/2})$, we have that the condition on $A_f A_g$ is satisfied. If in addition $\max_{i=1}^n |v_P(z_i)|$ and $\max_{i=1}^n |u_P(z_i)|$ are $O_P(\sqrt{n})$, then the conditions on B_f and B_g will be automatically satisfied. The latter conditions would hold if $\mathbb{E}_P u_P^2(Z) < \infty$ and $\mathbb{E}_P v_P^2(Z) < \infty$, for example.

Note that the rate of convergence requirement on A_f and A_g is a slower rate of convergence than the rate obtained when estimating Lipschitz regression functions when $d_Z = 1$, for example. Furthermore, we show in Section 4 that f and g being in a reproducing kernel Hilbert space (RKHS) is enough for them to be estimable at the required rate.

In the setting where Z is high-dimensional and f and g are sparse and linear, standard theory for the Lasso [10, 49] shows that it may be used to obtain estimates \hat{f} and \hat{g} satisfying the required properties under appropriate sparsity conditions. In fact, in this case our test statistic is closely related to that involved in the ANT procedure of Ren et al. [41] and the so-called RP test introduced in Shah and Bühlmann [46], which amount to a regularised partial correlation. A difference is that the denominator in (3) means the GCM test would not require $\varepsilon_P \perp \xi_P$ unlike the ANT test and the RP test.

We now briefly sketch the reason for the relatively weak requirement on the MSPEs. In the following, we suppress dependence on P for simplicity of presentation. We have

$$(5) \quad \begin{aligned} \frac{1}{\sqrt{n}} \sum_{i=1}^n R_i &= \frac{1}{\sqrt{n}} \sum_{i=1}^n \{f(z_i) - \hat{f}(z_i) + \varepsilon_i\} \{g(z_i) - \hat{g}(z_i) + \xi_i\} \\ &= (b + v_g + v_f) + \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i \xi_i, \end{aligned}$$

where

$$b := \frac{1}{\sqrt{n}} \sum_{i=1}^n \{f(z_i) - \hat{f}(z_i)\} \{g(z_i) - \hat{g}(z_i)\},$$

$$v_g := \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i \{g(z_i) - \hat{g}(z_i)\}, \quad v_f := \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \{f(z_i) - \hat{f}(z_i)\}.$$

The summands in the final term in (5) are i.i.d. with zero mean provided $P \in \mathcal{P}_0$, so the central limit theorem dictates that these converge to a standard normal. We also see that the simple form of the GCM gives rise to the term b involving a product of bias-type terms from estimating f and g , so each term is only required to converge to 0 at a slow rate such that their product is of smaller order than the variance of the final term. The summands in v_g are, under the null, mean zero conditional on $(\mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$. This term and similarly v_f are therefore both relatively well behaved, and give rise to the weak conditions on B_f and B_g .

3.1.1. *Power of the GCM.* We now present a result on the power of a version of the GCM. We may view our test statistic as a normalised version of the conditional covariance $\mathbb{E}_P(\varepsilon_P \xi_P) = \mathbb{E}_P \mathbf{cov}_P(X, Y|Z)$, where $\mathbf{cov}_P(X, Y|Z) = \mathbb{E}_P(XY|Z) - \mathbb{E}_P(X|Z)\mathbb{E}_P(Y|Z)$. This is always zero under the null; see equation (1), and does not necessarily need to be nonzero under an alternative; we can only hope to have power against alternatives where this conditional covariance is nonzero.

Control of the term b in (5) under the alternative can proceed in exactly the same way as under the null. However, control of the terms v_f and v_g typically requires additional conditions (e.g., Donsker-type conditions) on the estimators \hat{f} and \hat{g} as under the alternative both the errors ε_i and \hat{g} can depend on $\mathbf{Y}^{(n)}$. A notable exception is when f and g are sparse linear functions; in this setting, alternative arguments can be used to show the GCM with Lasso regressions has optimal power when Z has a sparse inverse covariance [41, 46].

To state a general result avoiding additional conditions, here we will suppose that \hat{f} and \hat{g} have been constructed from an auxiliary training sample, independent of the data $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$ (e.g., through sample-splitting); see, for example, [42, 52]. A drawback however, compared to the original GCM, is that the corresponding prediction error terms A_f and A_g are here out-of-sample prediction errors. These are typically more sensitive to the distribution of Z and larger than the in-sample prediction errors featuring in Theorem 6. For this reason, we consider the sample splitting approach to be more of a tool to facilitate theoretical analysis and would usually recommend using the original GCM in practice due to its typically better type I error control.

THEOREM 8. *Consider the setup of Theorem 6 but with the following differences: \hat{f} and \hat{g} have been constructed using auxiliary data independent of $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$; the null hypothesis \mathcal{P}_0 is replaced by \mathcal{E}_0 the set of all distributions absolutely continuous with respect to Lebesgue measure; and conditions involving $\varepsilon_P \xi_P$ are replaced by those involving the centred version $\varepsilon_P \xi_P - \mathbb{E}_P(\varepsilon_P \xi_P)$. Define*

$$\rho_P = \mathbb{E}_P \mathbf{cov}_P(X, Y|Z) \quad \text{and} \quad \sigma_P = \sqrt{\mathbf{var}_P(\varepsilon_P \xi_P)}.$$

Then under the conditions of (i) in Theorem 6 we have

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\frac{\tau_N^{(n)} - \sqrt{n} \rho_P}{\tau_D^{(n)}} \leq t \right) - \Phi(t) \right| \rightarrow 0, \quad \tau_D^{(n)} - \sigma_P = o_P(1),$$

with $\tau_N^{(n)}$ and $\tau_D^{(n)}$ defined as in (3).

Under the conditions of (ii) in Theorem 6, we have

$$\sup_{P \in \mathcal{P}} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\frac{\tau_N^{(n)} - \sqrt{n} \rho_P}{\tau_D^{(n)}} \leq t \right) - \Phi(t) \right| \rightarrow 0, \quad \tau_D^{(n)} - \sigma_P = o_P(1).$$

A proof is given in the supplementary material. We see that we achieve optimal \sqrt{n} rates for estimating ρ_P .

3.1.2. *Relationship to semiparametric models.* In viewing $\tau_N^{(n)} / \sqrt{n}$ as an estimator of the functional ρ_P , our GCM test connects to a vast literature in semiparametric statistics. In particular, the requirement of estimating nonparametric quantities (in our case $\sqrt{A_f}$ and $\sqrt{A_g}$) at a $o(n^{-1/4})$ rate is common for estimators of functionals based on estimating equations involving influence functions [9]. Our requirement on prediction error necessitates that at least one of f_P and g_P is Hölder β -smooth with $\beta / (2\beta + d_Z) \geq 1/4$. Estimators of the expected conditional covariance functional requiring minimal possible smoothness conditions may be derived using the theory of higher order influence functions [34, 42–44]; these estimators are however significantly more complicated. Newey and Robins [36] study another approach to estimation of the functional based on a particular spline-based regression method. The work of Chernozhukov et al. [14] uses related ideas to ours here to obtain $1/\sqrt{n}$ convergent estimates and confidence intervals for parameters such as average treatment effects in causal inference settings. A distinguishing feature of our work here is that we only require in-sample prediction error bounds under the null of conditional independence, which is advantageous in our setting for the reasons mentioned in the previous section.

3.2. *Multivariate X and Y.* We now consider the more general setting where $d_X, d_Y \geq 1$, and will assume for technical reasons that $d_X d_Y \geq 3$. We let $T_{jk}^{(n)}$ be the univariate GCM based on data $(\mathbf{X}_j^{(n)}, \mathbf{Y}_k^{(n)}, \mathbf{Z}^{(n)})$ and regression methods \hat{f}_j and \hat{g}_k . (As described in Section 1.4, the subindex selects a column.) More generally, we will add subscripts j and k to certain terms defined in the previous subsection to indicate that the quantities are based on X_j and Y_k rather than X and Y . Thus, for example, $\varepsilon_{P,j}$ is the difference of X_j and its conditional expectation given Z .

We define our aggregated test statistic to be

$$S_n = \max_{j=1, \dots, d_X, k=1, \dots, d_Y} |T_{jk}^{(n)}|.$$

There are other choices for how to combine the test statistics in $\mathbf{T}^{(n)} := (T_{jk}^{(n)})_{j,k} \in \mathbb{R}^{d_X \times d_Y}$ into a single test statistic. Under similar conditions to those in Theorem 6, one can show that if d_X and d_Y are fixed, $\mathbf{T}^{(n)}$ will converge in distribution to a multivariate Gaussian limit with a covariance that can be estimated. The continuous mapping theorem can then be used to deduce the asymptotic limit distribution of the sum of squares of $T_{jk}^{(n)}$, for example. However, one advantage of the maximum is that the bias component of S_n will be bounded by the maximum of the bias terms in $T_{jk}^{(n)}$. A sum of squares-type statistic would have a larger bias component, and tests based on it may not maintain the level for moderate to large d_X or d_Y . Furthermore, S_n will tend to exhibit good power against alternatives where conditional independence is only violated for a few pairs (X_j, Y_k) , that is, when the set of (j, k) such that $X_j \not\perp\!\!\!\perp Y_k | Z$ is small.

In order to understand what values of S_n indicate rejection, we will compare S_n to

$$\hat{S}_n = \max_{j=1, \dots, d_X, k=1, \dots, d_Y} |\hat{T}_{jk}^{(n)}|$$

where $\hat{\mathbf{T}}^{(n)} \in \mathbb{R}^{d_X \cdot d_Y}$ is mean zero multivariate Gaussian with a covariance matrix $\hat{\Sigma} \in \mathbb{R}^{d_X \cdot d_Y \times d_X \cdot d_Y}$ determined from the data as follows. Let $\mathbf{R}_{jk} \in \mathbb{R}^n$ be the vector of products of residuals (2) involved in constructing the test statistic $T_{jk}^{(n)}$. We set $\hat{\Sigma}_{jk,lm}$ to be the sample correlation between $\mathbf{R}_{jk} \in \mathbb{R}^n$ and \mathbf{R}_{lm} :

$$\hat{\Sigma}_{jk,lm} = \frac{\mathbf{R}_{jk}^T \mathbf{R}_{lm} - \bar{\mathbf{R}}_{jk} \bar{\mathbf{R}}_{lm}}{(\|\mathbf{R}_{jk}\|_2^2/n - \bar{\mathbf{R}}_{jk}^2)^{1/2} (\|\mathbf{R}_{lm}\|_2^2/n - \bar{\mathbf{R}}_{lm}^2)^{1/2}}.$$

Here, $\bar{\mathbf{R}}_{jk}$ is the sample mean of the components of \mathbf{R}_{jk} .

Let \hat{G} be the quantile function of \hat{S}_n . This is a random function that depends on the data $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$ through $\hat{\Sigma}$. Note that given the \mathbf{R}_{jk} , we can approximate \hat{G} to any degree of accuracy via Monte Carlo.

The ground-breaking work of Chernozhukov et al. [15] gives conditions under which \hat{G} can well approximate the quantile function of a version of S_n where all bias terms, that is, terms corresponding to b , v_g and v_f are all equal to 0. We will require that those conditions are met by $\varepsilon_{P,j} \xi_{P,k}$ for all $j = 1, \dots, d_X, k = 1, \dots, d_Y$. Below, we lay out these conditions, which take two possible forms. Let $d = \max(d_X, d_Y)$; note that d and \mathcal{P} are permitted to change with n , though we suppress this in the notation.

- (A1a) $\max_{r=1,2} \mathbb{E}_P(|\varepsilon_{P,j} \xi_{P,k}|^{2+r} / C_n^r) + \mathbb{E}_P(\exp(|\varepsilon_{P,j} \xi_{P,k}| / C_n)) \leq 4$;
- (A1b) $\max_{r=1,2} \mathbb{E}_P(|\varepsilon_{P,j} \xi_{P,k}|^{2+r} / C_n^{r/2}) + \mathbb{E}_P(\max_{j,k} |\varepsilon_{P,j} \xi_{P,k}|^4 / C_n^2) \leq 4$;
- (A2) $C_n^2 (\log(dn))^7 / n \leq Cn^{-c}$ and $\mathbb{E}_P(\varepsilon_{P,j}^2 \xi_{P,k}^2) \geq c_1$.

The result below shows that under the moment conditions above, provided the prediction error following the regressions goes to zero sufficiently fast, \hat{G} closely approximates the quantile function of S_n and, therefore, may be used to correctly calibrate our test.

THEOREM 9. *Suppose that for $\mathcal{P} \subset \mathcal{P}_0$, the following is true: there are constants $C, c, c_1 > 0$ such that for each n and $P \in \mathcal{P}$, there exists $C_n \geq 1$ such that one of (A1a) and (A1b) hold, and that (A2) holds. Suppose that*

$$(6) \quad \max_{j,k} A_{f,j} A_{g,k} = o_{\mathcal{P}}(n^{-1} \log(d)^{-2}).$$

Suppose further that there exist sequences $(\tau_{f,n})_{n \in \mathbb{N}}, (\tau_{g,n})_{n \in \mathbb{N}}$ such that

$$(7) \quad \max_{i,j} |\varepsilon_{P,ij}| = O_{\mathcal{P}}(\tau_{g,n}), \quad \max_k A_{g,k} = o_{\mathcal{P}}(\tau_{g,n}^{-2} \log(d)^{-4})$$

$$(8) \quad \max_{i,k} |\xi_{P,ik}| = O_{\mathcal{P}}(\tau_{f,n}), \quad \max_j A_{f,j} = o_{\mathcal{P}}(\tau_{f,n}^{-2} \log(d)^{-4}).$$

Then

$$\sup_{P \in \mathcal{P}} \sup_{\alpha \in (0,1)} |\mathbb{P}_P\{S_n \leq \hat{G}(\alpha)\} - \alpha| \rightarrow 0.$$

A proof is given in the supplementary material.

REMARK 10. If the errors $\{\varepsilon_{P,j}\}_{j=1}^{d_X}$ and $\{\xi_{P,k}\}_{k=1}^{d_Y}$ are all sub-Gaussian with parameters bounded above by some constant M uniformly across $P \in \mathcal{P}$, we may easily see that both (A1a) and (A1b) are satisfied with C_n a constant; see Chernozhukov et al. [15] for further discussion.

If additionally we have $A_{f,j}, A_{g,k} = o_{\mathcal{P}}(\log(d)^{-1} \min\{n^{-1/2}, \log(d)^{-4}\})$, (6), (7) and (8) will all be satisfied.

Theorem 9 allows for d_X and d_Y to be large compared to n . However, the result can be of use even when faced with univariate data. In this case, or more generally when d_X and d_Y are small, one can consider mappings $f_X : \mathbb{R}^{d_X+d_Z} \rightarrow \mathbb{R}^{\tilde{d}_X}$ and $f_Y : \mathbb{R}^{d_Y+d_Z} \rightarrow \mathbb{R}^{\tilde{d}_Y}$ where \tilde{d}_X and \tilde{d}_Y are potentially large. Provided these mappings are not determined from the data, we will have for $\tilde{X} := f_X(X, Z)$ and $\tilde{Y} := f_Y(Y, Z)$ that $\tilde{X} \perp\!\!\!\perp \tilde{Y} \mid Z$ if $X \perp\!\!\!\perp Y \mid Z$ (see equation (1)). Thus we may apply the methodology above to the mapped data, potentially allowing the test to have power against a more diverse set of alternatives. In view of Theorem 8, successful mappings should have the equivalent of ρ_P large, but also $\mathbb{E}(\tilde{X} \mid Z = \cdot)$ and $\mathbb{E}(\tilde{Y} \mid Z = \cdot)$ should not be so complex that it is impossible to estimate them well. We leave further investigation of this topic to further work.

4. GCM based on kernel ridge regression. We now apply the results of the previous section to a GCM based on estimating the conditional expectations via kernel ridge regression. For simplicity, we consider only the univariate case where $d_X = d_Y = 1$. In the following, we make use of the notation introduced in Section 3.1.

Given $\mathcal{P} \subset \mathcal{P}_0$, suppose that the conditional expectations f_P, g_P satisfy $f_P, g_P \in \mathcal{H}$ for some RKHS $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ with reproducing kernel $k : \mathbb{R}^{d_Z} \times \mathbb{R}^{d_Z} \rightarrow \mathbb{R}$. Let $K \in \mathbb{R}^{n \times n}$ have ij th entry $K_{ij} = k(z_i, z_j)/n$ and denote the eigenvalues of K by $\hat{\mu}_1 \geq \hat{\mu}_2 \geq \dots \geq \hat{\mu}_n \geq 0$. We will assume that under each $P \in \mathcal{P}$, k admits an eigen-expansion of the form

$$(9) \quad k(z, z') = \sum_{j=1}^{\infty} \mu_{P,j} e_{P,j}(z) e_{P,j}(z')$$

with orthonormal eigenfunctions $\{e_{P,j}\}_{j=1}^{\infty}$, so $\mathbb{E}_P e_{P,j} e_{P,k} = \mathbb{1}_{\{k=j\}}$, and summable eigenvalues $\mu_{P,1} \geq \mu_{P,2} \geq \dots \geq 0$. Such an expansion is guaranteed under mild conditions by Mercer’s theorem.

Consider forming estimates $\hat{f} = \hat{f}^{(n)}$ and $\hat{g} = \hat{g}^{(n)}$ through kernel ridge regressions of $\mathbf{X}^{(n)}$ and $\mathbf{Y}^{(n)}$ on $\mathbf{Z}^{(n)}$ in the following way. For $\lambda > 0$, let

$$\hat{f}_{\lambda} = \operatorname{argmin}_{h \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{i=1}^n \{x_i - h(z_i)\}^2 + \lambda \|h\|_{\mathcal{H}}^2 \right\}.$$

We will consider selecting a final tuning parameter $\hat{\lambda}$ in the following data-dependent way:

$$\hat{\lambda} = \operatorname{argmin}_{\lambda > 0} \left\{ \frac{1}{n} \sum_{i=1}^n \frac{\hat{\mu}_i^2}{(\hat{\mu}_i + \lambda)^2} + \lambda \right\}.$$

The term minimised on the RHS is an upper bound on the mean-squared prediction error omitting constant factors depending on σ^2 (defined below in Theorem 11) and $\|f_P\|_{\mathcal{H}}^2$ or $\|g_P\|_{\mathcal{H}}^2$. Because of the hidden dependence on these quantities, this is not necessarily a practically effective way of selecting λ : our use of it here is simply to facilitate theoretical analysis. Finally, define $\hat{f} = \hat{f}_{\hat{\lambda}}$, and define \hat{g} analogously. We will write $T^{(n)}$ for the test statistic formed as in (3) with these choices of \hat{f} and \hat{g} .

THEOREM 11. *Let \mathcal{P} be such that $u_P(z), v_P(z) \leq \sigma^2$ for all z and $P \in \mathcal{P}$.*

- (i) *For any $P \in \mathcal{P}$, $\sup_{t \in \mathbb{R}} |\mathbb{P}_P(T^{(n)} \leq t) - \Phi(t)| \rightarrow 0$.*
- (ii) *Suppose $\sup_{P \in \mathcal{P}} \mathbb{E}_P \{|\varepsilon_P \xi_P|^{2+\eta}\} \leq c$ for some $c \geq 0$ and $\eta > 0$. Suppose further that $\sup_{P \in \mathcal{P}} \max(\|f_P\|_{\mathcal{H}}, \|g_P\|_{\mathcal{H}}) < \infty$ and*

$$(10) \quad \limsup_{\lambda \downarrow 0} \sup_{P \in \mathcal{P}} \sum_{j=1}^{\infty} \min(\mu_{P,j}, \lambda) = 0.$$

Then

$$\sup_{P \in \mathcal{P}} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(T^{(n)} \leq t) - \Phi(t)| \rightarrow 0.$$

A proof is given the supplementary material.

REMARK 12. An application of dominated convergence theorem shows that a sufficient condition for (10) to hold is that $\sum_{j=1}^\infty \sup_{P \in \mathcal{P}} \mu_{P,j} < \infty$.

The proof proceeds by first showing that the ridge regression estimators \hat{f} and \hat{g} satisfy $A_f A_g = o_P(n^{-1})$ and then applies Theorem 6. The requirement that f_P and g_P lie in a RKHS satisfying (9) is a rather weak regularity condition on the conditional expectations. For example, taking the first-order Sobolev kernel shows that it is enough that the conditional expectations are Lipschitz when $d_Z = 1$, $\mathbb{P}(Z \in [0, 1]) = 1$ and the marginal density of Z is bounded above [1]. However, the uniformity offered by (ii) above requires $L := \sup_{P \in \mathcal{P}} \max(\|f_P\|_{\mathcal{H}}, \|g_P\|_{\mathcal{H}}) < \infty$ and a large value of L will require a large sample size in order for $T^{(n)}$ to have a distribution close to a standard normal. We investigate this, and evaluate the empirical performance of the GCM in the next section.

5. Experiments. Section 3 proposes the generalised covariance measure (GCM). Although we provide detailed computations for kernel ridge regression in Section 4, the technique can be combined with any regression method. In practice, the choice may depend on external knowledge of the specific application the user has in mind. In this section, we study the empirical performance of the GCM with boosted regression trees as the regression method. In particular, we use the R package `xgboost` [12, 13] with a ten-fold cross-validation scheme over the parameter `maxdepth`.

5.1. *No-free-lunch in conditional independence testing.* Theorem 2 states that if a conditional independence test has power against an alternative at a given sample size, then there is a distribution from the null that is rejected with probability larger than the significance level. We now illustrate the no-free-lunch theorem empirically.

Let us fix a RKHS \mathcal{H} that corresponds to a Gaussian kernel with bandwidth $\sigma = 1$. We now compute for different sample sizes the rejection rates for data sets generated from the following model: $Z = N_Z$, $Y = f_a(Z) + N_Y$, and $X = f_a(Z) + N_X$, with $N_X, N_Y, N_Z \sim \mathcal{N}(0, 1)$, i.i.d., and $f_a(z) := \exp(-z^2/2) \sin(az)$ defining a function $f_a \in \mathcal{H}$. Figure 2 shows a plot of f_a for $a = 6$ and $a = 18$. Clearly, for any a , we have $X \perp\!\!\!\perp Y \mid Z$, but for large values of a the independence will be harder to detect from data. We now fix three different sample sizes $n = 100$, $n = 1000$ and $n = 10,000$. For any of such sample size n , we can find an a , that is, a distribution from the null, such that the probability of (falsely) rejecting $X \perp\!\!\!\perp Y \mid Z$ is larger than the prespecified level α . Figure 3 shows the results for the GCM test with boosted regression trees and the significance level $\alpha = 0.05$: for any sample size, there

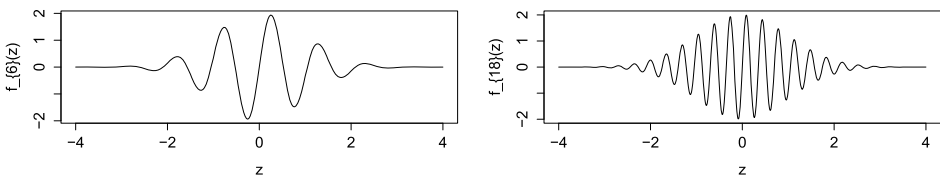


FIG. 2. Graphs of the function f_a for $a = 6$ (left) and $a = 18$ (right). This function is used as the conditional mean that needs to be estimated from data. The RKHS norm increases exponentially with a ; see (11).

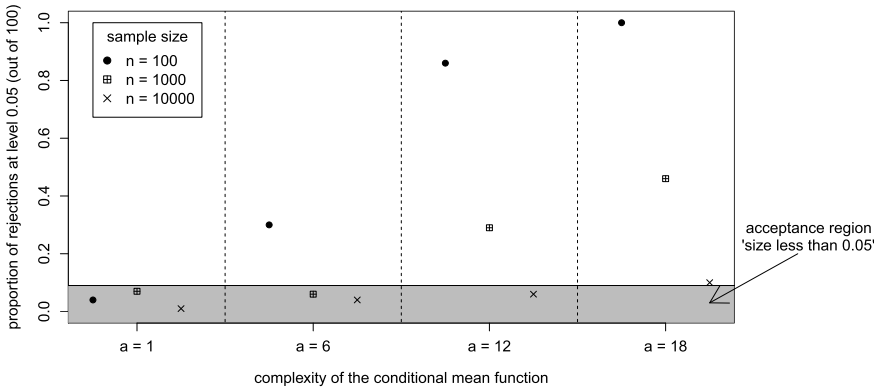


FIG. 3. Illustration of the no-free-lunch theorem; see Section 5.1. No sample size is large enough to ensure the correct level for all distributions from the null: there is always a distribution from the null which yields a type I error that is larger than the pre-specified significance level of 0.05. The shaded area indicates the area in which we accept the null hypothesis that the size of the test is less than 0.05.

exists a distribution from the null, for which the test rejects the null hypothesis of conditional independence. For $n = 100$, we can choose $a = 6$, for $n = 1000$, we choose $a = 12$, and for $n = 10,000$, $a = 18$.

This sequence of distributions violates one of the assumptions that we require for the GCM test to obtain uniform asymptotic level guarantee. Intuitively, for large a , the conditional expectations $z \mapsto \mathbb{E}[X|Z = z]$ and $z \mapsto \mathbb{E}[Y|Z = z]$ are too complex to be estimated reliably from the data. More formally, the RKHS norm of the functions f_a are defined as

$$(11) \quad \|f_a\|_{\mathcal{H}}^2 = \int_{-\infty}^{\infty} F_a(\omega)^2 \exp(\sigma^2 \omega^2 / 2) d\omega = \sqrt{8\pi} \cdot (\exp(a^2) + \exp(-a^2)),$$

where

$$F_a(\omega) = \exp(-(\omega - a)^2 / 2) + \exp(-(\omega + a)^2 / 2)$$

is the Fourier transform of f_a . Equation (11) shows that a null hypothesis \mathcal{P} containing all of the above models for $a > 0$, violates one of the assumptions in Theorem 11: for this choice of RKHS and null hypothesis, there is no M such that $\sup_{P \in \mathcal{P}} \max(\|f_P\|_{\mathcal{H}}, \|g_P\|_{\mathcal{H}}) < M$. (Note that not all sequences of functions with growing RKHS norm also yield a violation of level guarantees: some functions with large RKHS norm, for example, modifications of constant functions, can be easily learned from data.) Other conditional independence tests fail on the examples in Figure 3, too, for a similar reason. However, most of these other methods are less transparent in the underlying assumptions, since they do not come with uniform level guarantees.

5.2. *On level and power.* It is of course impossible to provide an exhaustive simulation-based level and power size analysis. We therefore concentrate on a small choice of distributions from the null and the alternative. In the following, we compare the GCM with three other conditional independence tests: KCI [51] with its implementation from `CondIndTests` [26], and the residual prediction test [26, 46]. We also compare to a test that performs the same regression as GCM, but then tests for independence between the residuals, rather than vanishing correlation, using HSIC [25]. (This procedure is similar to the one that Fan et al. [20] propose to use in the case of additive noise models.) As we discuss in Example 1, we do not expect this test to hold level in general. We then consider the following distributions from the null:

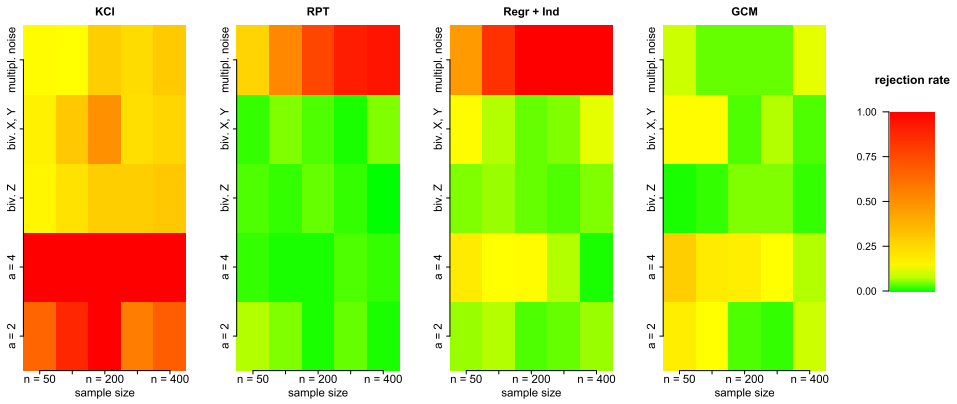


FIG. 4. *Level analysis: the GCM can hold the level if the sample size is large enough to reliably estimate the conditional mean. Testing for independence between residuals does not hold the level (third plot).*

- (a) $Z \sim \mathcal{N}(0, 1)$, $X = f_a(Z) + 0.3\mathcal{N}(0, 1)$, $Y = f_a(Z) + 0.3\mathcal{N}(0, 1)$, $a = 2$;
- (b) the same as (a) but with $a = 4$;
- (c) $Z_1, Z_2 \sim \mathcal{N}(0, 1)$ independent, $X = f_1(Z_1) - f_1(Z_2) + 0.3\mathcal{N}(0, 1)$, $Y = f_1(Z_1) + f_1(Z_2) + 0.3\mathcal{N}(0, 1)$;
- (d) $Z \sim \mathcal{N}(0, 1)$, $X_1 = f_1(Z) + 0.3\mathcal{N}(0, 1)$, $X_2 = f_1(Z) + X_1 + 0.3\mathcal{N}(0, 1)$, $Y_1 = f_1(Z) + 0.3\mathcal{N}(0, 1)$, $Y_2 = f_1(Z) + Y_1 + 0.3\mathcal{N}(0, 1)$; and
- (e) $Z \sim \mathcal{N}(0, 1)$, $Y = f_2(Z) \cdot \mathcal{N}(0, 1)$, $X = f_2(Z) \cdot \mathcal{N}(0, 1)$.

In the remainder of this section, we refer to these settings as (a) “ $a = 2$ ”, (b) “ $a = 4$ ”, (c) “biv. Z ”, (d) “biv. X, Y ”, and (e) “multipl. noise”, respectively. For each of the sample sizes 50, 100, 200, 300 and 400, we first generate 100 data sets, and then compute rejection rates of the considered conditional independence tests. The results are shown in Figure 4. For rejection rates below 0.11, the hypothesis “the size of the test is less than 0.05” is not rejected at level 0.01 (pointwise). The GCM indeed has promising behaviour in terms of type I error control. As expected, however, it requires the sample size to be big enough to obtain a reliable estimate for the conditional mean.

We then investigate the tests’ power by altering the data generating processes (a)–(e), described above. Each equation for Y receives an additional term $+0.2X$, which yields $X \not\perp\!\!\!\perp Y \mid Z$ (for (d), we add the term $+0.2X_2$ to the equation of Y_2). Figure 5 shows empirical rejection rates. All methods, except for RPT, are able to correctly reject the hypothesis

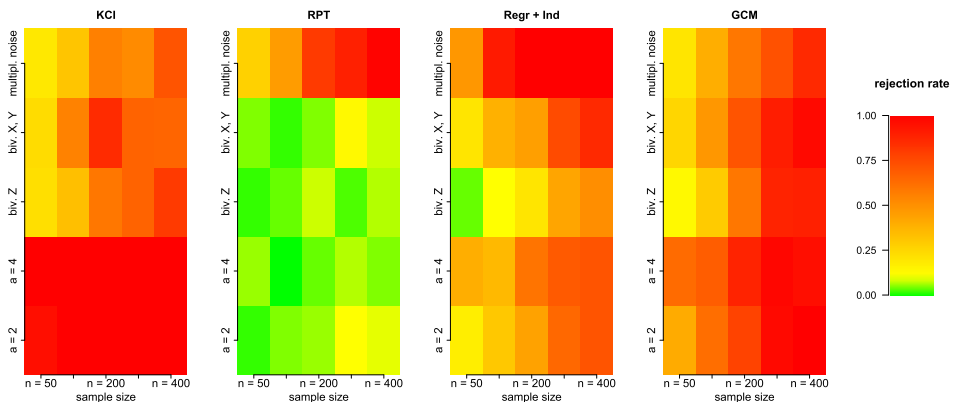


FIG. 5. *Power analysis: most of the methods are able to detect if the distribution does not satisfy conditional independence, in particular if the sample size increases.*

that the distribution is from the null, particularly with increasing sample size. In our experimental setup, it is the level analysis that poses a greater challenge for the methods other than GCM.

6. Discussion. A key result of this paper is that conditional independence testing is hard: nontrivial tests that maintain valid level over the entire class of distributions satisfying conditional independence and that are absolutely continuous with respect to Lebesgue measure cannot exist. In unconditional independence testing, control of type I error is straightforward and research efforts have focussed on power properties of tests. Our result indicates that in conditional independence testing, the basic requirement of type I error control deserves further attention. We argue that as domain knowledge is necessary in order to select a conditional independence test appropriate for a particular setting, there is a need to develop conditional independence tests whose suitability is reasonably straightforward to judge.

In this work, we have introduced the GCM framework to address this need. The ability for the GCM to maintain the correct level relies almost exclusively on the predictive properties of the regression procedures upon which it is based. Selecting a good regression procedure, whilst mathematically an equally impossible problem, can at least be usefully informed by domain knowledge. We hope to see further applications of GCM-based tests in the future. On the theoretical side, it would be interesting to understand more precisely the tradeoff between the type I and type II errors in conditional independence testing. Often, work on testing fixes a null and then considers what sorts of classes of alternative distributions it is possible, or impossible to maintain power against. In the context of conditional independence testing, the problem set is even richer, in that one must also consider subclasses of null distributions, and can then study power properties associated with that null.

APPENDIX A: PROOF OF THEOREM 2

The proof of Theorem 2 relies heavily on Lemma 13 in Section A.2, which shows that given any distribution Q where $(X, Y, Z) \sim Q$, one can construct $(\tilde{X}, \tilde{Y}, \tilde{Z})$ with $\tilde{X} \perp\!\!\!\perp \tilde{Y} \mid \tilde{Z}$ where $(\tilde{X}, \tilde{Y}, \tilde{Z})$ and (X, Y, Z) are arbitrarily close in ℓ_∞ -norm with arbitrarily high probability.

In the proofs of Theorem 2 and Lemma 13 below, we often suppress dependence on n to simplify the presentation. Thus, for example, we write \mathbf{X} for $\mathbf{X}^{(n)}$. We use the following notation. We write $s = (d_X + d_Y + d_Z)$ and will denote by $V \in \mathbb{R}^s$ the triple (X, Y, Z) . Furthermore, $\mathbf{V} := (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$. We denote by $p_{X,Y,Z}$ the density of (X, Y, Z) with respect to Lebesgue measure. We will use μ to denote Lebesgue measure on \mathbb{R}^{ns+1} and write Δ for the symmetric difference operator.

A.1. Proof of Theorem 2. Suppose, for a contradiction, that there exists a Q with support strictly contained in a ℓ_∞ -ball of radius M under which $X \not\perp\!\!\!\perp Y \mid Z$ but $\mathbb{P}_Q(\psi_n(\mathbf{V}; U) = 1) = \beta > \alpha$. We will henceforth assume that $V \sim Q$ and $\mathbf{V} := (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ are i.i.d. copies of V . Thus we may omit the subscript Q applied to probabilities and expectations in the sequel. Denote the rejection region by

$$R = \{(\mathbf{x}, \mathbf{y}, \mathbf{z}; u) \in \mathbb{R}^{ns} \times [0, 1] : \psi_n(\mathbf{x}, \mathbf{y}, \mathbf{z}; u) = 1\}.$$

Our proof strategy is as follows. Using Lemma 13, we will create $\tilde{V} := (\tilde{X}, \tilde{Y}, \tilde{Z})$ such that $\tilde{X} \perp\!\!\!\perp \tilde{Y} \mid \tilde{Z}$ but \tilde{V} is suitably close to V such that a corresponding i.i.d. sample $\tilde{\mathbf{V}} := (\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}) \in \mathbb{R}^{ns}$ satisfies $\mathbb{P}((\tilde{\mathbf{V}}, U) \in R) > \alpha$, contradicting that ψ_n has valid level α . How close \tilde{V} needs to be to V in order for this argument to work depends on the rejection region R . As an arbitrary Borel subset of $\mathbb{R}^{ns} \times [0, 1]$, R can be arbitrarily complex. In order to get a

handle on it, we will construct an approximate version R^\sharp of R that is a finite union of boxes; see Lemma 15.

Let $\eta = (\beta - \alpha)/7 > 0$. Since $\{(x, y, z) : p_{X,Y,Z}(x, y, z) > m\} =: B_m \downarrow \emptyset$ as $m \uparrow \infty$, there exists M_1 such that $\mathbb{P}((X, Y, Z) \in B_{M_1}^c) > 1 - \eta/n$. Let Ω_1 be the event that $(x_i, y_i, z_i) \in B_{M_1}^c$ for all $i = 1, \dots, n$. (Here and below, an event refers to an element in the underlying σ -algebra. Recall that x_i, y_i , and z_i denote rows of \mathbf{X}, \mathbf{Y} , and \mathbf{Z} , respectively, that is, they are random vectors.) Then by a union bound we have $\mathbb{P}(\Omega_1) \geq 1 - \eta$.

Let M_2 be such that $\mathbb{P}(\|\mathbf{V}\|_\infty > M_2) < \eta$ and let Ω_2 be the event that $\|\mathbf{V}\|_\infty \leq M_2$. Further define

$$\check{R} = \{(\mathbf{x}, \mathbf{y}, \mathbf{z}, u) \in R : \|(\mathbf{x}, \mathbf{y}, \mathbf{z})\|_\infty \leq M_2\}.$$

Note that

$$(12) \quad \mathbb{P}((\mathbf{V}, U) \in \check{R}) \geq \beta - \mathbb{P}((\mathbf{V}, U) \in R \setminus \check{R}) > \beta - \eta.$$

Let $L = L(\eta)$ be as defined in Lemma 13 (taking $\delta = \eta$). From Lemma 15 applied to \check{R} , we know there exists a finite union R^\sharp of hypercubes each of the form

$$\prod_{k=1, \dots, ns+1} (a_k, b_k]$$

such that $\mu(R^\sharp \Delta \check{R}) < \eta / \max(L, M_1^n)$. Now on the region $B_{M_1}^c$ defining Ω_1 we know that the density of (\mathbf{V}, U) is bounded above by M_1^n . Thus we have that

$$(13) \quad \mathbb{P}(\{(\mathbf{V}, U) \in \check{R} \setminus R^\sharp\} \cap \Omega_1) < \eta.$$

Now for $r \geq 0$ and $\mathbf{v} \in \mathbb{R}^{ns+1}$ let $B_r(\mathbf{v}) \subset \mathbb{R}^{ns+1}$ denote the ℓ_∞ ball with radius $r > 0$ and center \mathbf{v} . Define

$$R^r = \{\mathbf{v} \in R : B_r(\mathbf{v}) \subseteq R^\sharp\}.$$

Then since $R^r \uparrow R^\sharp$ as $r \downarrow 0$, there exists $r_0 > 0$ such that $\mu(R^\sharp \setminus R^{r_0}) < \eta / M_1^n$.

For $\epsilon = r_0$ and $B = R^\sharp \setminus \check{R}$, the statement of Lemma 13 provides us with $\tilde{\mathbf{V}} := (\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}})$ which satisfies $\mathbb{P}((\tilde{\mathbf{V}}, U) \in B) < L\mu(B) < \eta$ and with which we argue as follows. Let Ω_3 be the event that $\|\mathbf{V} - \tilde{\mathbf{V}}\|_\infty < r_0$, so $\mathbb{P}(\Omega_3) \geq 1 - \eta$.

$$\begin{aligned} \mathbb{P}((\tilde{\mathbf{V}}, U) \in R) &\geq \mathbb{P}((\tilde{\mathbf{V}}, U) \in \check{R}) \geq \mathbb{P}((\tilde{\mathbf{V}}, U) \in R^\sharp) - \mathbb{P}((\tilde{\mathbf{V}}, U) \in R^\sharp \setminus \check{R}) \\ &> \mathbb{P}(\{(\tilde{\mathbf{V}}, U) \in R^\sharp\} \cap \Omega_3) - \eta > \mathbb{P}((\mathbf{V}, U) \in R^{r_0}) - 2\eta \\ &\geq \mathbb{P}((\mathbf{V}, U) \in R^\sharp) - \mathbb{P}(\{(\mathbf{V}, U) \in R^\sharp \setminus R^{r_0}\} \cap \Omega_1) - \mathbb{P}(\Omega_1^c) - 2\eta \\ &> \mathbb{P}((\mathbf{V}, U) \in R^\sharp) - 4\eta. \end{aligned}$$

Now

$$\begin{aligned} \mathbb{P}((\mathbf{V}, U) \in R^\sharp) &\geq \mathbb{P}((\mathbf{V}, U) \in \check{R}) - \mathbb{P}(\{(\mathbf{V}, U) \in \check{R} \setminus R^\sharp\} \cap \Omega_1) - \mathbb{P}(\Omega_1^c) \\ &> \mathbb{P}((\mathbf{V}, U) \in \check{R}) - 2\eta > \beta - 3\eta \end{aligned}$$

using (13) and (12). Putting things together, we have $\mathbb{P}((\tilde{\mathbf{V}}, U) \in R) > \beta - 7\eta > \alpha$, completing the proof.

A.2. Auxiliary lemmas.

LEMMA 13. *Let (X, Y, Z) have a $(d_X + d_Y + d_Z)$ -dimensional distribution in $\mathcal{Q}_{0,M}$ for some $M \in (0, \infty]$. Let $(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)})$ be a sample of n i.i.d. copies of (X, Y, Z) . Given $\delta > 0$, there exists $L = L(\delta)$ such that for all $\epsilon > 0$ and all Borel subsets $B \subseteq \mathbb{R}^{n \cdot (d_X + d_Y + d_Z)} \times [0, 1]$, it is possible to construct n i.i.d. random vectors $(\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{Z}}^{(n)})$ with distribution $P \in \mathcal{P}_{0,M}$ where the following properties hold:*

- (i) $\mathbb{P}(\|(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}, \mathbf{Z}^{(n)}) - (\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{Z}}^{(n)})\|_\infty < \epsilon) > 1 - \delta$;
- (ii) *If $U \sim U[0, 1]$ independently of $(\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{Z}}^{(n)})$ then*

$$\mathbb{P}((\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{Z}}^{(n)}, U) \in B) \leq L\mu(B).$$

PROOF. We will first describe the construction of $\tilde{V} := (\tilde{X}, \tilde{Y}, \tilde{Z})$ from $V := (X, Y, Z)$. The corresponding n -sample $\tilde{\mathbf{V}} := (\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{Z}}^{(n)})$ will have observation vectors formed in the same way from the corresponding observation vectors in \mathbf{V} . The proof proceeds in three steps. We begin by creating a bounded version $\check{V} = (\check{X}, \check{Y}, \check{Z})$ of V supported on a grid $2^{-r}\mathbb{Z}$, for which we can control an upper bound on the probability mass function. Next, we apply Lemma 14 to obtain transforms $\check{V}^{(1)}, \dots, \check{V}^{(K!)}$ of \check{V} for arbitrarily large K where \check{X} has been ‘embedded’ in the last component. Then we create noisy versions $\{\check{V}^{(m)}\}_{m=1}^{K!}$ by adding uniform noise such that truncation of their binary expansions yields the discrete versions $\{\check{V}^{(m)}\}_{m=1}^{K!}$. Each of these are potential candidates for the random vector \tilde{V} , but we must ensure that the corresponding n -fold product obeys (ii). This is problematic as the embedding procedure necessarily creates near-degenerate random vectors that fall within small regions with large probability. To overcome this issue, we employ in the final step, a probabilistic argument that exploits the property, supplied by Lemma 14, that the $K!$ embeddings have supports with little overlap.

Step 1: Define $s := d_X + d_Y + d_Z$. Since $\{(x, y, z) \in \mathbb{R}^s : p_{X,Y,Z}(x, y, z) > t\} =: B_t \downarrow \emptyset$ as $t \uparrow \infty$, there exists M_1 such that the event $\Lambda_1 = \{(X, Y, Z) \in B_{M_1}^c\}$ has $\mathbb{P}(\Lambda_1) \geq 1 - \delta/(2n)$. Next, let $M_2 < M$ be such that $\mathbb{P}(\|V\|_\infty > M_2) < \delta/(2n)$, and let Λ_2 be the event that $\|V\|_\infty \leq M_2$. For later use, we define the events

$$\Omega_1 = \{(x_i, y_i, z_i) \in B_{M_1}^c \text{ for all } i = 1, \dots, n\} \quad \text{and} \quad \Omega_2 = \{\|V\|_\infty \leq M_2\}.$$

Note that union bounds give $\mathbb{P}((\Omega_1 \cap \Omega_2)^c) < \delta$.

Let $E^{(1)}$ be uniformly distributed on $[-M_2, M_2]^s$. Let $r \in \mathbb{N}$ be such that $2^{-r} < \min(\epsilon/3, (M - M_2)/3, 1/n)$ and define

$$\check{V} := (\check{X}, \check{Y}, \check{Z}) := 2^{-r} \lfloor 2^r (V \mathbb{1}_{\Lambda_1 \cap \Lambda_2} + E^{(1)} \mathbb{1}_{(\Lambda_1 \cap \Lambda_2)^c}) \rfloor.$$

Here, the floor function is applied componentwise. Note that \check{V} takes values in a grid $(2^{-r}\mathbb{Z})^s$ and satisfies

$$(14) \quad \|\check{V} - V\|_\infty \mathbb{1}_{\Lambda_1 \cap \Lambda_2} \leq 2^{-r} < \epsilon/3.$$

The choice of r ensures that $\check{V} \in (-M', M')^s$ where $M' = M - 2(M - M_2)/3$. Furthermore, the inclusion of the $\mathbb{1}_{\Lambda_1 \cap \Lambda_2}$ term and $E^{(1)}$ ensures that the probability it takes any given value is bounded above by $M_3 2^{-sr}$ where $M_3 := \max(M_1, (M_2/2)^{-s})$ is independent of ϵ . Indeed for any fixed $k \in \mathbb{Z}^s$, writing $A = [k2^{-r}, (k+1)2^{-r})$ we have

$$\mathbb{P}(V \in A | \Lambda_1 \cap \Lambda_2) \leq M_1 2^{-rs} \quad \text{and} \quad \mathbb{P}(E^{(1)} \in A | (\Lambda_1 \cap \Lambda_2)^c) = 2^{-rs} / (2M_2)^s.$$

As $\mathbb{P}(\check{V} = k2^{-r})$ is a convex combination of these probabilities, it must be at most their maximum.

Step 2: We can now apply Lemma 14 with $W = (\check{Y}, \check{Z})$ and $N = \check{X}$. This gives us $K!$ random vectors $\check{V}^{(1)}, \dots, \check{V}^{(K!)}$ where $K > 2^r > n$; for each $m = 1, \dots, K!$, $\check{V}^{(m)} = (\check{X}^{(m)}, \check{Y}^{(m)}, \check{Z}^{(m)})$ satisfies

- (a) $\mathbb{P}(|\mathring{V}_s^{(m)} - \check{V}_s| \leq 2^{-r}) = 1$ and $\mathring{V}_j^{(m)} = \check{V}_j$ for $j \leq s - 1$;
- (b) $\mathring{X}^{(m)}$ may be recovered from $\mathring{Z}^{(m)}$ via $\mathring{X}^{(m)} = \mathring{g}_m(\mathring{Z}^{(m)})$ for some function \mathring{g}_m ;
- (c) $\mathring{V}_s^{(m)}$ takes values in $K^{-2}2^{-r}\mathbb{Z}$ and the probability it takes any given value is bounded above by $2^{-sr}K^{-1}M_3$;

and, additionally,

(d) the supports $\mathring{S}_1, \dots, \mathring{S}_{K!}$ of $\mathring{V}^{(1)}, \dots, \mathring{V}^{(K!)}$ obey the following structure: there exists a collection of K^2 disjoint sets $\{\mathring{G}_{jk}\}_{j,k=1}^K$ and an enumeration $\pi_1, \dots, \pi_{K!}$ of the permutations of $\{1, \dots, K\}$ such that $\mathring{S}_m = \bigcup_{k=1}^K \mathring{G}_{k\pi_m(k)}$ and $\mathbb{P}(\mathring{V}^{(m)} \in \mathring{G}_{j\pi_m(j)}) = K^{-1}$ for all $m = 1, \dots, K!$ and $j = 1, \dots, K$.

We now create a noisy version of the $\mathring{V}^{(m)}$ that obeys similar properties to the above, but is absolutely continuous with respect to Lebesgue measure. To this end, we introduce $E^{(2)} = (E_X, E_Y, E_Z) \in (0, 1)^s$ with independent $U(0, 1)$ components. Then let $\tilde{V}^{(m)} \in \mathbb{R}^s$ be defined by

$$\tilde{V}_j^{(m)} = \begin{cases} K^{-2}2^{-r}E_s^{(2)} + \mathring{V}_s^{(m)} & \text{for } j = s, \\ 2^{-r}E_j^{(2)} + \mathring{V}_j^{(m)} & \text{otherwise.} \end{cases}$$

This obeys

- (a') $\mathbb{P}(\|\tilde{V}^{(m)} - \mathring{V}^{(m)}\|_\infty \leq 2^{-r}) = 1$;
- (b') $\tilde{X}^{(m)}$ may be recovered from $\tilde{Z}^{(m)}$ via $\tilde{X}^{(m)} = g_k(\tilde{Z}^{(m)})$ for some function g_m , which depends on K and r ;
- (c') the density of $\tilde{V}^{(m)}$ with respect to Lebesgue measure is bounded above by KM_3 (indeed, using (c), we have that it is bounded by $2^{rs}K^2 \cdot 2^{-sr}K^{-1}M_3 = KM_3$);
- (d') the supports $S_1, \dots, S_{K!}$ of the $\{\tilde{V}^{(m)}\}_{m=1}^{K!}$ obey property (d) with the disjoint sets \mathring{G}_{jk} above replaced by the Minkowski sum $G_{jk} = \mathring{G}_{jk} + 2^{-r}((0, 1)^{s-1} \times (0, K^{-2}))$.

Note that (b') holds as we can first construct $\mathring{Z}^{(m)}$ from $\tilde{Z}^{(m)}$ and then apply (b). The former is done by removing the additive noise component by truncating the binary expansion appropriately: $\mathring{Z}^{(m)} := K^{-2}2^{-r} \lfloor K^2 2^r \tilde{Z}^{(m)} \rfloor$. A consequence of this property is that decomposing $(\tilde{X}^{(m)}, \tilde{Y}^{(m)}, \tilde{Z}^{(m)}) = \tilde{V}^{(m)}$, we have $\tilde{X}^{(m)} \perp\!\!\!\perp \tilde{Y}^{(m)} \mid \tilde{Z}^{(m)}$. To see this, we argue as follows. Let us write p_A and $p_{A|B}$ for the densities of A and A given B , respectively, when A and B are random vectors. Suppressing dependence on m temporarily, we have that for any \tilde{z} with $p_{\tilde{Z}}(\tilde{z}) > 0$,

$$\begin{aligned} p_{\tilde{X}, \tilde{Y}, \tilde{Z}}(\tilde{x}, \tilde{y} \mid \tilde{z}) &= p_{E_X, \tilde{Y} \mid \tilde{Z}}(\tilde{x} - g(\tilde{z}), \tilde{y} \mid \tilde{z}) = p_{E_X}(\tilde{x} - g(\tilde{z})) p_{\tilde{Y} \mid \tilde{Z}}(\tilde{y} \mid \tilde{z}) \\ &= p_{\tilde{X} \mid \tilde{Z}}(\tilde{x} \mid \tilde{z}) p_{\tilde{Y} \mid \tilde{Z}}(\tilde{y} \mid \tilde{z}), \end{aligned}$$

so $\tilde{X} \perp\!\!\!\perp \tilde{Y} \mid \tilde{Z}$.

Property (d') follows as the support of each $\mathring{V}^{(m)}$ is contained in $2^{-r}(\mathbb{Z}^{s-1} \times K^{-2}\mathbb{Z})$ (see (a) and (c)).

From (a'), by the triangle inequality we have that

$$(15) \quad \mathbb{P}(\|(\tilde{V}^{(m)} - V)\mathbb{1}_{\{\Lambda_1 \cap \Lambda_2\}}\|_\infty \leq \epsilon) = 1$$

and $\tilde{V}^{(m)} \in (-M, M)^s$. Let $\{\tilde{\mathbf{V}}^{(m)}\}_{m=1}^{K!}$ be the corresponding n -sample versions of $\{\tilde{V}^{(m)}\}_{m=1}^{K!}$. Then for any m , (15) gives $\mathbb{P}(\|(\tilde{\mathbf{V}}^{(m)} - \mathbf{V})\mathbb{1}_{\Omega_1 \cap \Omega_2}\|_\infty \leq \epsilon) = 1$. Thus $\mathbb{P}(\|\tilde{\mathbf{V}}^{(m)} - \mathbf{V}\|_\infty \leq \epsilon) > 1 - \delta$. We see that any $\tilde{\mathbf{V}}^{(m)}$ satisfies all requirements of the result except potentially (ii).

Step 3: In order to pick an m for which (ii) is satisfied, we use the so-called probabilistic method. First, note we may assume $0 < \mu(B) < \infty$ or otherwise any m will do. Define

$T_m := S_m^n \times [0, 1]$ to be the support set of $(\tilde{V}^{(m)}, U)$ where $U \sim U[0, 1]$ independently of $\{\tilde{V}^{(m)}\}_{m=1}^{K!}$.

For $j \in \{1, \dots, K\}$, let $G_j := \bigcup_{k=1}^K G_{jk}$. Let \mathcal{J} be the set of n -tuples (j_1, \dots, j_n) of distinct elements of $\{1, \dots, K\}$. Now define

$$(16) \quad C := \bigcup_{(j_1, \dots, j_n) \in \mathcal{J}} \prod_{l=1}^n G_{j_l}$$

and let $D = C \times [0, 1]$. Fix $(j_l)_{l=1}^n \in \mathcal{J}$ and $(k_l)_{l=1}^n \in \{1, \dots, K\}^n$, and set $G := \prod_{l=1}^n G_{j_l k_l}$. Then G has nonempty intersection with a given S_m^n if and only if $k_l = \pi_m(j_l)$ for all l . Thus if $(k_l)_{l=1}^n \notin \mathcal{J}$, G is disjoint from all S_m^n . On the other hand if $(k_l)_{l=1}^n \in \mathcal{J}$, the number of S_m^n that intersect G is $(K - n)!$, the number of permutations of $\{1, \dots, K\}$ whose outputs are fixed at n points. We therefore have that all but at most $(K - n)!$ of the support sets T_m are disjoint from $G \times [0, 1]$, whence

$$\sum_{m=1}^{K!} \mu\{(G \times [0, 1]) \cap B \cap T_m\} \leq (K - n)! \mu\{(G \times [0, 1]) \cap B\}.$$

Now the set C is the disjoint union of all sets $\prod_{l=1}^n G_{j_l k_l}$ with $(j_l)_{l=1}^n \in \mathcal{J}$ and $(k_l)_{l=1}^n \in \{1, \dots, K\}^n$. Thus summing over all such sets we obtain

$$\sum_{m=1}^{K!} \mu(D \cap B \cap T_m) \leq (K - n)! \mu(D \cap B) \leq (K - n)! \mu(B).$$

This gives that there exists at least one $m = m^*$ with

$$\mu(B \cap D \cap T_{m^*}) \leq \frac{(K - n)!}{K!} \mu(B).$$

Next, observe that the number of cells $\prod_{l=1}^n G_{k_l}$ where at least two of k_1, \dots, k_n are the same is $K^n - K(K - 1) \dots (K - n + 1)$. As $\mathbb{P}(\tilde{V}^{(m)} \in G_j) = K^{-1}$ for all j , using (16) we have

$$\mathbb{P}((\tilde{V}^{(m)}, U) \notin D) = K^{-n} \{K^n - K(K - 1) \dots (K - n + 1)\} = O(K^{-1})$$

for every m . Putting things together, we have that there must exist a m^* with

$$\begin{aligned} \mathbb{P}((\tilde{V}^{(m^*)}, U) \in B) &\leq \mathbb{P}((\tilde{V}^{(m^*)}, U) \in B \cap D) + \mathbb{P}((\tilde{V}^{(m^*)}, U) \notin D) \\ &\leq K^n M_3^n \frac{(K - n)!}{K!} \mu(B) + O(K^{-1}) \leq 2M_3^n \mu(B) \end{aligned}$$

for K sufficiently large, which can be arranged by taking r sufficiently large. \square

LEMMA 14. *Let $W \in \mathbb{R}^l$, $N \in \mathbb{R}^d$ be random vectors. Suppose that N is bounded and there is some $r \in \mathbb{N}$ such that both W and N have components taking values in the grid $2^{-r}\mathbb{Z}$. Suppose further that the probability that (N, W) takes any particular value is bounded by $2^{-(m+d)r} M$ for some $M > 0$. Then there exists a $K \in \mathbb{N}$ with $K > 2^r$ and a set of $K!$ functions $\{f_1, \dots, f_{K!}\}$ where*

$$(W_l, N) \mapsto f_m(W_l, N) =: \hat{W}_l^{(m)} \in \mathbb{R}$$

such that for each $m = 1, \dots, K!$:

- (i) $\mathbb{P}(|W_l - \hat{W}_l^{(m)}| \leq 2^{-r}) = 1$;
- (ii) there is some function $g_m : \mathbb{R} \rightarrow \mathbb{R}^d$ such that $N = g_m(\hat{W}_l^{(m)})$;
- (iii) $\hat{W}_l^{(m)}$ has components taking values in a grid $K^{-2}2^{-r}\mathbb{Z}$.

Moreover, defining $\mathring{W}^{(m)} = (W_1, \dots, W_{l-1}, \mathring{W}_l^{(m)})$,

(iv) the probability that $(N, \mathring{W}^{(m)})$ takes any value is bounded above by $K^{-1}2^{-(m+d)r}M$;

(v) the supports $S_1, \dots, S_{K!}$ of $(N, \mathring{W}^{(1)}), \dots, (N, \mathring{W}^{(K!)})$ obey the following structure: there exists K^2 disjoint sets $\{G_{jk}\}_{j,k=1}^K$ and an enumeration $\pi_1, \dots, \pi_{K!}$ of the permutations of $\{1, \dots, K\}$ such that for all m , $S_m = \bigcup_{k=1}^K G_{k\pi_m(k)}$, and for all m and k , $\mathbb{P}((N, \mathring{W}^{(m)}) \in G_{k\pi_m(k)}) = K^{-1}$.

PROOF. As N is bounded, by replacing $g_m(\cdot)$ by $g_m(\cdot) + v$ where $v \in \mathbb{R}^d$ is appropriately chosen with components in $2^{-r}\mathbb{Z}$, we may assume that N has nonnegative components. Let $t \in \mathbb{N}$ be such that $2^t > 2^r \max(1, \|N\|_\infty)$. We shall prove the result with $K = 2^{dt}$.

Define the random variable \mathring{N} by

$$\mathring{N} = 2^r \sum_{j=0}^{d-1} 2^{tj} N_{j+1}.$$

This is a concatenation of the binary expansions of $2^r N_j \in \{0, 1, 2, \dots, 2^t - 1\}$ for $j = 1, \dots, d$. Observe that $\mathring{N} \in \{0, 1, \dots, K - 1\}$ and that N_j may be recovered from \mathring{N} by examining its binary expansion. Indeed, $2^r N_j$ is the residue modulo 2^t of $\lfloor \mathring{N} / 2^{t(j-1)} \rfloor$.

For $j = \{0, 1, \dots, K - 1\}$, let \tilde{N}_j be the residue of $\mathring{N} + j$ modulo K , so $\tilde{N}_j \in \{0, 1, \dots, K - 1\}$. Also, for $k = \{0, 1, \dots, K - 1\}$ let $\tilde{N}_{j,k} = \tilde{N}_j + Kk$. Note that $\tilde{N}_{j,k}$ takes values in $\{0, \dots, K^2 - 1\}$. Let the random variable E be uniformly distributed on $\{0, 1, \dots, K - 1\}$ independently of all other quantities. Now let $\pi_1, \dots, \pi_{K!}$ be an enumeration of the permutations of $\{0, \dots, K - 1\}$. Finally, let $\mathring{N}_m = \tilde{N}_{E, \pi_m(E)}$ for $m = 1, \dots, K!$.

One important feature of this construction is that we can recover E from \mathring{N}_m (and m) via $\pi_m(E) = \lfloor \mathring{N}_m / K \rfloor$, and thereby determine E , which then reveals \mathring{N} and each of the individual N_j . In summary, this gives us $K!$ different embeddings of the vector N into a single random variable.

We may now define f_m by

$$f_m(W_l, N) = W_l + 2^{-r} K^{-2} \mathring{N}_m.$$

It is easy to see that (i) and (iii) are satisfied. To deduce (ii), observe that we may recover W_l via $2^{-r} \lfloor 2^r f_m(W_l, N) \rfloor =: c_m(f_m(W_l, N))$, and thus also determine \mathring{N}_m which, as discussed above, also gives us N and E . Let g_m and h_m be the functions that when applied to $f_m(W_l, N)$, yield N and E , respectively. Let us introduce the notation that for a vector $v \in \mathbb{R}^s$, $v_{-j} \in \mathbb{R}^{s-1}$ for $j = 1, \dots, s$ is the subvector of v where the j th component is omitted. Then we have

$$\begin{aligned} \mathbb{P}(N = n, \mathring{W}^{(m)} = \mathring{w}) &= \mathbb{P}(W_{-l} = \mathring{w}_{-l}, W_l = c_m(\mathring{w}_l), N = n, E = h_m(\mathring{w}_l)) \mathbb{1}_{\{g_m(\mathring{w}_l) = n\}} \\ &= K^{-1} \mathbb{P}(W_{-l} = \mathring{w}_{-l}, W_l = c_m(\mathring{w}_l), N = n) \mathbb{1}_{\{g_m(\mathring{w}_l) = n\}} \\ &\leq K^{-1} 2^{-(l+d)r} M, \end{aligned}$$

using the independence of E in the second line above. This gives (iv).

Note that the supports of the $(N, \tilde{N}_{j,k})$ are all disjoint as (j, k) can be recovered from $(N, \tilde{N}_{j,k})$. For $j, k = 0, 1, \dots, K - 1$ let G_{jk} be the support set of

$$(N, \tilde{W}_{j,k}) := (N, W_1, \dots, W_{l-1}, W_l + 2^{-r} K^{-2} \tilde{N}_{j,k}).$$

From the above, we see that the $\{G_{jk}\}_{j,k=0}^{K-1}$ are all disjoint. Property (v) follows from noting that $\mathring{W}^{(m)} = \tilde{W}_{E, \pi_m(E)}$. \square

The following well-known result appears for example in Weaver [50], Theorem 2.19.

LEMMA 15. *Given any bounded Borel subset B of \mathbb{R}^d and any $\epsilon > 0$, there exists a finite union of boxes of the form*

$$B^\sharp = \bigcup_{i=1}^N \prod_{k=1}^d (a_{i,k}, b_{i,k}]$$

such that $\mu(B \Delta B^\sharp) \leq \epsilon$, where μ denotes Lebesgue measure and Δ denotes the symmetric difference operator.

APPENDIX B: KL-SEPARATION OF NULL AND ALTERNATIVE

For distributions $P_1, P_2 \in \mathcal{E}_0$, let $\text{KL}(P_1 \parallel P_2)$ denote the KL divergence from P_2 to P_1 , which we define to be $+\infty$ when P_1 is not absolutely continuous with respect to P_2 . The following result, which is proved in the supplementary material, shows that it is possible to choose $Q \in \mathcal{Q}_0$ that is arbitrarily far away from \mathcal{P}_0 (by picking $\sigma^2 > 0$ to be sufficiently small).

PROPOSITION 16. *Consider the distribution Q over the triple $(X, Y, Z) \in \mathbb{R}^3$ defined in the following way: $Y = X + N$ with $X \sim \mathcal{N}(0, 1)$, $N \sim \mathcal{N}(0, \sigma^2)$ and $X \perp\!\!\!\perp N$. The variable $Z \sim \mathcal{N}(0, 1)$ is independent of (X, Y) . Thus $X \not\perp\!\!\!\perp Y \mid Z$, that is, $Q \in \mathcal{Q}_0$, and we have*

$$\inf_{P \in \mathcal{P}_0} \text{KL}(Q \parallel P) = \frac{1}{2} \log\left(\frac{1 + \sigma^2}{\sigma^2}\right) = \inf_{P \in \mathcal{P}_0} \text{KL}(P \parallel Q).$$

Acknowledgements. We thank Kacper Chwialkowski, Kenji Fukumizu, Arthur Gretton, Bernhard Schölkopf and Ilya Tolstikhin for helpful discussions, initiated by BS, on the hardness of conditional independence testing, and KC and AG for helpful comments on the manuscript. BS has raised the point that for finitely many data conditional independence testing may be arbitrarily hard in several of his talks, for example, at the Machine Learning Summer School in Tübingen in 2013. We thank Matey Neykov and Cyrill Scheidegger for kindly pointing out errors in earlier versions of this manuscript, and suggesting potential fixes. We also thank Peter Bühlmann for helpful discussions regarding the aggregation of tests via taking the maximum test statistic. Finally, we thank four anonymous referees and an Associate Editor for helpful comments that have improved the manuscript.

The first author was supported by The Alan Turing Institute under the EPSRC Grant EP/N510129/1, an EPSRC Programme Grant and an EPSRC First Grant.

The second author was supported by a research grant (18968) from VILLUM FONDEN and the Alan Turing Institute supported JP's research visit in London (July 2017).

SUPPLEMENTARY MATERIAL

Supplement to “The hardness of conditional independence testing and the generalised covariance measure” (DOI: [10.1214/19-AOS1857SUPP](https://doi.org/10.1214/19-AOS1857SUPP); .pdf). This contains the proofs of some of the main results as well as a few auxiliary statements.

REFERENCES

- [1] BACH, F. (2017). On the equivalence between kernel quadrature rules and random feature expansions. *J. Mach. Learn. Res.* **18** Paper No. 21, 38. [MR3634888](https://arxiv.org/abs/1706.02515)

- [2] BAHADUR, R. R. and SAVAGE, L. J. (1956). The nonexistence of certain statistical procedures in non-parametric problems. *Ann. Math. Stat.* **27** 1115–1122. MR0084241 <https://doi.org/10.1214/aoms/1177728077>
- [3] BALAKRISHNAN, S. and WASSERMAN, L. (2019). Hypothesis testing for densities and high-dimensional multinomials: Sharp local minimax rates. *Ann. Statist.* **47** 1893–1927. MR3953439 <https://doi.org/10.1214/18-AOS1729>
- [4] BARRON, A. R. (1989). Uniformly powerful goodness of fit tests. *Ann. Statist.* **17** 107–124. MR0981439 <https://doi.org/10.1214/aos/1176347005>
- [5] BERGSMAN, W. P. (2004). Testing conditional independence for continuous random variables. EURANDOM-report 2004-049.
- [6] BERRETT, T. B. and SAMWORTH, R. J. (2019). Nonparametric independence testing via mutual information. *Biometrika* **106** 547–566. MR3992389 <https://doi.org/10.1093/biomet/asz024>
- [7] BERRETT, T. B., WANG, Y., BARBER, R. F. and SAMWORTH, R. J. (2018). The conditional permutation test. ArXiv, arxiv:1807.05405.
- [8] BERTANHA, M. and MOREIRA, M. J. (2018). Impossible inference in econometrics: Theory and applications. ArXiv, arxiv:1612.02024v3.
- [9] BICKEL, P. J., KLAASSEN, C. A. J., RITOV, Y. and WELLNER, J. A. (1993). *Efficient and Adaptive Estimation for Semiparametric Models*. Johns Hopkins Series in the Mathematical Sciences. Johns Hopkins Univ. Press, Baltimore, MD. MR1245941
- [10] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Springer, Heidelberg. MR2807761 <https://doi.org/10.1007/978-3-642-20192-9>
- [11] CANAY, I. A., SANTOS, A. and SHAIKH, A. M. (2013). On the testability of identification in some nonparametric models with endogeneity. *Econometrica* **81** 2535–2559. MR3138554 <https://doi.org/10.3982/ECTA10851>
- [12] CHEN, T. and GUESTRIN, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SigKDD International Conference on Knowledge Discovery and Data Mining* 785–794. ACM.
- [13] CHEN, T., HE, T., BENESTY, M., KHOTILOVICH, V. and TANG, Y. (2018). xgboost: Extreme gradient boosting. R package version 0.6.4.1.
- [14] CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., DUFLO, E., HANSEN, C., NEWEY, W. and ROBINS, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **21** C1–C68. MR3769544 <https://doi.org/10.1111/ectj.12097>
- [15] CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *Ann. Statist.* **41** 2786–2819. MR3161448 <https://doi.org/10.1214/13-AOS1161>
- [16] DAUDIN, J.-J. (1980). Partial association measures and an application to qualitative regression. *Biometrika* **67** 581–590. MR0601095 <https://doi.org/10.1093/biomet/67.3.581>
- [17] DAWID, A. P. (1979). Conditional independence in statistical theory. *J. Roy. Statist. Soc. Ser. B* **41** 1–31. MR0535541
- [18] DOBRA, A., HANS, C., JONES, B., NEVINS, J. R., YAO, G. and WEST, M. (2004). Sparse graphical models for exploring gene expression data. *J. Multivariate Anal.* **90** 196–212. MR2064941 <https://doi.org/10.1016/j.jmva.2004.02.009>
- [19] DUFOUR, J.-M. (2003). Identification, weak instruments, and statistical inference in econometrics. *Can. J. Econ.* **36** 767–808.
- [20] FAN, J., FENG, Y. and XIA, L. (2015). A projection based conditional dependence measure with applications to high-dimensional undirected graphical models. ArXiv, arxiv:1501.01617.
- [21] FISHER, R. A. (1920). A mathematical examination of the methods of determining the accuracy of an observation by the mean error and by the mean square error. *Mon. Not. R. Astron. Soc.* **80** 758–770.
- [22] FISHER, R. A. (1934). Two new properties of mathematical likelihood. *Proc. R. Soc. A* **114** 285–307.
- [23] FISHER, R. A. (1935). *The Design of Experiments*. Oliver and Boyd, Edinburgh, UK.
- [24] FUKUMIZU, K., GRETTON, A., SUN, X. and SCHÖLKOPF, B. (2008). Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems 20 (NIPS)* 489–496. MIT Press.
- [25] GRETTON, A., FUKUMIZU, K., TEO, C. H., SONG, L., SCHÖLKOPF, B. and SMOLA, A. (2008). A kernel statistical test of independence. In *Advances in Neural Information Processing Systems 20 (NIPS)* 585–592. MIT Press.
- [26] HEINZE-DEML, C., PETERS, J. and MEINSHAUSEN, N. (2017). Invariant causal prediction for nonlinear models. *J. Causal Inference*. To appear. ArXiv, arxiv:1706.08576.
- [27] Hoeffding, W. (1948). A non-parametric test of independence. *Ann. Math. Stat.* **19** 546–557. MR0029139 <https://doi.org/10.1214/aoms/1177730150>

- [28] JENSEN, J. L. and SØRENSEN, M. (2017). *Statistical principles: A first course (lecture notes)*, Aarhus and Copenhagen, Denmark, available upon request.
- [29] KIM, S. (2015). ppcor: Partial and semi-partial (part) correlation. R package version 1.1.
- [30] KOLLER, D. and FRIEDMAN, N. (2009). *Probabilistic Graphical Models: Principles and Techniques. Adaptive Computation and Machine Learning*. MIT Press, Cambridge, MA. [MR2778120](#)
- [31] KRAFT, C. (1955). Some conditions for consistency and uniform consistency of statistical procedures. *Univ. California Publ. Statist.* **2** 125–141. [MR0073896](#)
- [32] LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series 17*. Oxford University Press, New York. [MR1419991](#)
- [33] LECAM, L. (1973). Convergence of estimates under dimensionality restrictions. *Ann. Statist.* **1** 38–53. [MR0334381](#)
- [34] LI, L., TCHETGEN TCHETGEN, E., VAN DER VAART, A. and ROBINS, J. M. (2011). Higher order inference on a treatment effect under low regularity conditions. *Statist. Probab. Lett.* **81** 821–828. [MR2793749](#) <https://doi.org/10.1016/j.spl.2011.02.030>
- [35] MARKOWETZ, F. and SPANG, R. (2007). Inferring cellular networks—a review. *BMC Bioinform.* **8 Suppl 6** S5. <https://doi.org/10.1186/1471-2105-8-S6-S5>
- [36] NEWBY, W. K. and ROBINS, J. (2018). Cross-fitting and fast remainder rates for semiparametric estimation. ArXiv, [arxiv:1801.09138](https://arxiv.org/abs/1801.09138).
- [37] PEARL, J. (2009). *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge Univ. Press, Cambridge. [MR2548166](#) <https://doi.org/10.1017/CBO9780511803161>
- [38] PETERS, J., BÜHLMANN, P. and MEINSHAUSEN, N. (2016). Causal inference by using invariant prediction: Identification and confidence intervals. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 947–1012. With comments and a rejoinder. [MR3557186](#) <https://doi.org/10.1111/rssb.12167>
- [39] PETERS, J., JANZING, D. and SCHÖLKOPF, B. (2017). *Elements of Causal Inference: Foundations and Learning Algorithms. Adaptive Computation and Machine Learning*. MIT Press, Cambridge, MA. [MR3822088](#)
- [40] RAMSEY, J. D. (2014). A scalable conditional Independence test for nonlinear, non-Gaussian data. ArXiv, [arxiv:1401.5031](https://arxiv.org/abs/1401.5031).
- [41] REN, Z., SUN, T., ZHANG, C.-H. and ZHOU, H. H. (2015). Asymptotic normality and optimalities in estimation of large Gaussian graphical models. *Ann. Statist.* **43** 991–1026. [MR3346695](#) <https://doi.org/10.1214/14-AOS1286>
- [42] ROBINS, J., LI, L., TCHETGEN TCHETGEN, E. and VAN DER VAART, A. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman* 335–421. IMS.
- [43] ROBINS, J., TCHETGEN TCHETGEN, E., LI, L. and VAN DER VAART, A. (2009). Semiparametric minimax rates. *Electron. J. Stat.* **3** 1305–1321. [MR2566189](#) <https://doi.org/10.1214/09-EJS479>
- [44] ROBINS, J. M., LI, L., MUKHERJEE, R., TCHETGEN, E. T. and VAN DER VAART, A. (2017). Minimax estimation of a functional on a structured high-dimensional model. *Ann. Statist.* **45** 1951–1987. [MR3718158](#) <https://doi.org/10.1214/16-AOS1515>
- [45] ROMANO, J. P. (2004). On non-parametric testing, the uniform behaviour of the t -test, and related problems. *Scand. J. Stat.* **31** 567–584. [MR2101540](#) <https://doi.org/10.1111/j.1467-9469.2004.00407.x>
- [46] SHAH, R. D. and BÜHLMANN, P. (2018). Goodness-of-fit tests for high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 113–135. [MR3744714](#) <https://doi.org/10.1111/rssb.12234>
- [47] SHAH, R. D. and PETERS, J. (2020). Supplement to “The hardness of conditional independence testing and the generalised covariance measure.” <https://doi.org/10.1214/19-AOS1857SUPP>.
- [48] SPIRITES, P., GLYMOUR, C. and SCHEINES, R. (2000). *Causation, Prediction, and Search*, 2nd ed. *Adaptive Computation and Machine Learning*. MIT Press, Cambridge, MA. With additional material by David Heckerman, Christopher Meek, Gregory F. Cooper and Thomas Richardson, A Bradford Book. [MR1815675](#)
- [49] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- [50] WEAVER, N. (2013). *Measure Theory and Functional Analysis*. World Scientific Co. Pte. Ltd., Hackensack, NJ. [MR3185963](#) <https://doi.org/10.1142/8813>
- [51] ZHANG, K., PETERS, J., JANZING, D. and SCHÖLKOPF, B. (2011). Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the 27th Annual Conference on Uncertainty in Artificial Intelligence (UAI)* 804–813. AUAI Press.
- [52] ZHENG, W. and VAN DER LAAN, M. J. (2011). Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning. Springer Ser. Statist.* 459–474. Springer, New York. [MR2867139](#) https://doi.org/10.1007/978-1-4419-9782-1_27