



The Hellinger correlation

Gery Geenens & Pierre Lafaye de Micheaux

To cite this article: Gery Geenens & Pierre Lafaye de Micheaux (2020): The Hellinger correlation, Journal of the American Statistical Association, DOI: [10.1080/01621459.2020.1791132](https://doi.org/10.1080/01621459.2020.1791132)

To link to this article: <https://doi.org/10.1080/01621459.2020.1791132>



View supplementary material [↗](#)



Accepted author version posted online: 09 Jul 2020.



Submit your article to this journal [↗](#)



Article views: 124



View related articles [↗](#)



View Crossmark data [↗](#)

The Hellinger correlation

Gery Geenens*, Pierre Lafaye de Micheaux

School of Mathematics and Statistics, UNSW Sydney, Australia

*Corresponding author: ggeenens@unsw.edu.au, School of Mathematics and Statistics, UNSW Australia, Sydney, tel +61 2 938 57032, fax +61 2 9385 7123

Abstract

In this paper, the defining properties of any valid measure of the dependence between two continuous random variables are revisited and complemented with two original ones, shown to imply other usual postulates. While other popular choices are proved to violate some of these requirements, a class of dependence measures satisfying all of them is identified. One particular measure, that we call the Hellinger correlation, appears as a natural choice within that class due to both its theoretical and intuitive appeal. A simple and efficient nonparametric estimator for that quantity is proposed, with its implementation publicly available in the R package `HellCor`. Synthetic and real-data examples illustrate the descriptive ability of the measure, which can also be used as test statistic for exact independence testing.

Keywords: Dependence, Measure, Copula, Hellinger distance, Nearest-Neighbor

1 Introduction

A large part of science is about understanding the dependence between several factors that may influence each other, for instance to disentangle genetics and environmental risk factors for individual diseases. Hence statistics, the art of turning empirical evidence into information, has always kept dependence modelling at its core. Characterising the dependence between two variables includes two main tasks: (i) testing the null hypothesis H_0 that the two variables are independent; and (ii) measuring the strength of any dependence that may exist between the two variables.

These two tasks are often amalgamated, for instance when a dependence measure is used as the test statistic in an independence testing procedure, or when the observed value of the test statistic is interpreted as a quantification of the underlying dependence. This may seem natural, however there are good reasons to approach them separately. Indeed a measure is expected to be a fair quantification of the strength of the involved dependence. In a case of weak but non-null dependence, we would expect a reliable measure of

dependence to take a value accordingly low. By contrast, independence testing aims at making a binary decision as to the presence of dependence. As such, a powerful test should be based on a statistic that takes values as large as possible (i.e., very different than under H_0) as soon as dependence is present, regardless of its strength. Consequently, using an interpretable dependence measure as a test statistic might result in some loss of power for the resulting test, while a test statistic designed for guaranteeing high power for the resulting test might lack a bit of finesse for accurately quantifying the strength of the dependence. We illustrate the former point in Section 5.5 through a small simulation study. The latter can be understood through the observation made in Section 3.2, that some popular ‘measures’ may take arbitrarily large values even in some situations of ‘near-independence’, e.g. when the dependence is *very* low but of specific type. In such scenarios, those quantities, used as independence test statistics, would certainly achieve very high power. However, they do not *measure* accurately the real level of dependence involved. See also [Reimherr and Nicolae \(2013\)](#), [Sun and Zhao \(2014\)](#) for related comments. This paper focuses on meaningfully measuring dependence, without explicitly giving a central role to concepts inspired by dichotomous testing procedures, such as power.

The literature on quantifying dependence has long been monopolised by the historical measures, such as Pearson’s correlation, Spearman’s ρ , Kendall’s τ and Hoeffding’s D . Yet, the interest in modern alternatives has recently made an impressive upsurge. Among others, one can cite distance correlation ([Székely et al, 2007](#)), maximal information criterion ([Reshef et al, 2011](#)), Hilbert-Schmidt independence criterion ([Gretton et al, 2005](#), [Pfister et al, 2018](#)), sign covariance ([Bergsma and Dassios, 2014](#)), G-squared ([Wang et al, 2017](#)) and symmetric rank covariance ([Weihs et al, 2018](#)), along with a renewed enthusiasm for the mutual information ([Kinney and Atwal, 2014](#), [Zeng et al, 2018](#), [Berrett and Samworth, 2019](#)).

Through this abundance resurfaces the question of the criteria discriminating sensible measures from others. Far from new, it was already addressed by [Rényi \(1959\)](#), who formulated 7 axioms that a valid dependence measure

between two variables, say X_1 and X_2 , should satisfy. However, the only measure known to fulfil them all is the *maximal correlation coefficient*, viz. $S(X_1, X_2) = \sup_{\psi_1, \psi_2} |\rho(\psi_1(X_1), \psi_2(X_2))|$, where $\rho(\cdot, \cdot)$ is Pearson's correlation and the supremum is taken over all Borel functions $\psi_1, \psi_2 : \mathbb{R} \rightarrow \mathbb{R}$. Yet, that measure is computable only in very particular cases, and it may return, and often does so, its maximal value $S(X_1, X_2) = 1$ even in the absence of any obvious strong dependence between X_1 and X_2 (Bell, 1962). This evidences that some of Renyi's axioms may be unsuitable for general use. As a result, different sets of amended axioms have been proposed in the subsequent literature, see e.g., Schweizer and Wolff (1981), Lancaster (1982), Granger et al (2004) or Balakrishnan and Lai (2009, Section 4.3). Among those propositions, five properties, labelled (P1)–(P5) hereafter, seem difficult to dispute, while others are more prone to discussion. Below, we complement those 5 mainstream postulates with two original ones (P6)–(P7), and justify at length their reason-of-being. We show that they are more fundamental than other usually posited properties, while more natural intuitively.

2 Renyi's axioms and beyond

2.1 Dependence

We call *dependence* between two random variables X_1 and X_2 whatever remains to be specified for entirely reconstructing the joint distribution F_{12} of (X_1, X_2) once their marginal distributions F_1 and F_2 are known. The strength of dependence is thus the size (in an appropriate sense) of that missing link. As such, X_1 and X_2 are as dependent as can be when F_{12} is as different as can be to the independence base case $F_1 F_2$. If both X_1 and X_2 are continuous, this is characterised by F_{12} being singular with respect to $F_1 F_2$ (Silvey, 1964). In the discrete case, though, such singularity is impossible. This illustrates why measuring dependence may be a structurally different problem in the continuous and in the discrete cases (Hoeffding, 1942). In particular, it is known that approaches based on copulas, warranted in the continuous setting, are doomed to failure for analysing dependence between discrete variables (Genest and Neslehová, 2007). This justifies to study the two

situations separately; this paper considers the continuous case only. Perspectives of extension of the below discussion to the discrete setting are provided in Section 7.

2.2 Postulates

Let X_1 and X_2 be two continuous random variables defined on the same probability space. It is widely accepted that a valid measure D of the dependence between them should be such that:

- (P1) (*existence*) $D(X_1, X_2)$ is defined, whatever the variables X_1 and X_2 ;
- (P2) (*symmetry*) $D(X_1, X_2) = D(X_2, X_1)$;
- (P3) (*normalisation*) $0 \leq D(X_1, X_2) \leq 1$;
- (P4) (*characterisation of independence*) $D(X_1, X_2) = 0 \Leftrightarrow X_1$ and X_2 are independent ($X_1 \perp\!\!\!\perp X_2$);
- (P5) (*weak Gaussian conformity*) If (X_1, X_2) is a bivariate Gaussian vector, then $D(X_1, X_2)$ is a strictly increasing function of $|\rho(X_1, X_2)|$.

‘Existence’ (P1) is a minimal requirement. Though, many popular measures do not satisfy it. For instance, Pearson’s correlation and distance correlation (dCor) are only defined if X_1 and X_2 have finite second moment.

Defined as the void between F_{12} and $F_1 F_2$ (Section 2.1), the dependence in (X_1, X_2) is evidently invariant to permutation of X_1 and X_2 , making ‘symmetry’ (P2) unquestionable as well. Note that asymmetric measures, arguments in favour of which may sometimes be found in the literature, explicitly target directional notions such as ‘regression dependence’ (Dette et al, 2013) or causal relationships (Janzing et al, 2013), hence are of a different nature and do not measure dependence as defined hereinabove.

‘Normalisation’ (P3) only aims to provide benchmarks – any candidate measure can be made comply with it through renormalisation. Importantly, though, it implies that $D(X_1, X_2)$ is an unsigned number. Any signed measure,

whose sign is meant to be informative about the ‘direction’ of the association ($D(X_1, X_2) > 0$: positive association, X_1 and X_2 tend to vary in the same direction; $D(X_1, X_2) < 0$: negative association, X_1 and X_2 tend to vary in opposite direction) is meaningful only when explicitly looking for such monotonic association. Dependence, as defined in Section 2.1, is a much broader concept and cannot generally be categorised as ‘positive’ or ‘negative’. For instance, among the scatter-plots shown in Figure 5.2 below, none (but [7]) exhibits any sense of ‘positive’ or ‘negative’ association between X_1 and X_2 while all (but [6]) involve dependence between them. Hence a general dependence measure must be unsigned. Then it seems only natural to ask the measure to be null in the case and only in the case of no dependence (P4).

Finally, ‘weak Gaussian conformity’ (P5) is unavoidable. In a bivariate Gaussian vector, Pearson’s correlation ρ uniquely specifies the joint distribution once the marginals are fixed. Hence dependence (as defined in Section 2.1) is unequivocally characterised by Pearson’s correlation, and any measure disagreeing with it in a bivariate Gaussian vector cannot be valid.

Mostly dictated by common sense, these (P1)–(P5) can be found under this form or slightly amended in most of the related references. Here, we complete this list with the following two original requirements which, by contrast, offer novel perspectives on what characterises valid dependence measures.

(P6) (*characterisation of pure dependence*) $D(X_1, X_2) = 1$ if and only if there exists a Borel function $\Psi: [0,1] \rightarrow \mathbb{R}^2$ such that $(X_1, X_2) = \Psi(U)$ for $U \sim U_{[0,1]}$ and C_Ψ , the image of Ψ in \mathbb{R}^2 , is such that $\iint_{C_\Psi} dF_1(x_1)dF_2(x_2) = 0$;

(P7) (*generalised Data Processing Inequality*) If X_1 and X_3 are conditionally independent given X_2 ($X_1 \perp\!\!\!\perp X_3 \mid X_2$), then $D(X_1, X_3) \leq \min\{D(X_1, X_2), D(X_2, X_3)\}$.

Analogously to (P4), $D(X_1, X_2)$ should be maximum if and only if there exists some sort of perfect dependence between X_1 and X_2 . Yet, a universally

accepted definition of what perfect dependence is, has proved elusive. Our interpretation of it, which we refer to as *pure dependence* and leads to (P6), aligns closely with Hoeffding (1942)'s and Silvey (1964)'s conception as described in Section 2.3. The rationale for (P7) is detailed in Section 2.4 and is shown to have wide implications.

2.3 Pure dependence versus predictability

One can view a vector (X_1, X_2) whose components are independent as a random system with two degrees of freedom, in the sense that X_1 and X_2 are allowed to vary totally freely. By contrast, any degree of dependence between X_1 and X_2 necessarily restrains, to some extent, their free variation, reducing *de facto* the associated number of degrees of freedom for (X_1, X_2) to strictly smaller than 2. This number can be thought of as the (possibly fractional) number of latent variables able to reproduce in principle the joint behaviour of (X_1, X_2) . From that perspective, the opposite of 'independence' is thus when (X_1, X_2) has only one degree of freedom, that is, when one single latent variable, say U , is able to generate the full covariation of X_1 and X_2 . Formally, this means that there is a function $\Psi: [0,1] \rightarrow \mathbb{R}^2$ such that $(X_1, X_2) = \Psi(U)$. Although we are to see two variables X_1 and X_2 , the underlying probabilistic process is fed by one single source of variability, and X_1 and X_2 are just the two sides of the same coin. This is essentially what we refer to as *pure dependence*. The concept is strongly related to Hoeffding (1942)'s '*c*-dependence', and is akin to the joint distribution F_{12} being singular with respect to the product $F_1 F_2$ of its margins (Silvey, 1964), although not exactly equivalent (Durante et al, 2013).

Indeed, the condition $\iint_{C_\Psi} dF_1(x_1)dF_2(x_2) = 0$ in (P6) implies that C_Ψ must be a 'proper curve', in the specific sense that the intersection of C_Ψ with any line $x_1 = t$ or $x_2 = t, t \in \mathbb{R}$, consists almost surely of at most finitely many points. In particular, it excludes the situations where, although seeded by one single random variable U , $(X_1, X_2) = \Psi(U)$ defines a couple of independent random variables. Clearly, this is the case when Ψ defines a constant function of either X_1 or X_2 , making X_1 or X_2 a degenerate variable hence independent of

any other. It is also the case, for instance, when Ψ is a *space-filling curve* such as the Peano or the Hilbert curve. Those are known to be surjective and continuous functions from the unit interval $I \subseteq [0,1]$ onto the unit square I^2 ; hence, as u varies from 0 to 1, $\Psi(u)$ visits every single point of I^2 (see Figure 5.3 below). In addition, they have the bi-measure-preserving property (Steele, 1997, Section 2.7): for any Borel set $A \subseteq I$, $\lambda_1(A) = \lambda_2(\Psi(A))$, where for $q = 1, 2$, λ_q is the Lebesgue measure on \mathbb{R}^q . In theory, it is thus possible to generate a bivariate uniform vector (X_1, X_2) on the unit square, hence $X_1 \perp\!\!\!\perp X_2$, by letting U run on I and defining $(X_1, X_2) = \Psi(U)$ (Steele, 1997, p. 43). Of course, by definition the image of this function is the whole I^2 , and $\iint_{I^2} dF_1(x_1)dF_2(x_2) = 1$ which violates our definition of pure dependence (P6).

Those space-filling curves are special cases of fractal constructions, and the observed independence of X_1 and X_2 is actually induced by the inherent chaos in the fractal Ψ . A fractal is obtained as the limit of a series of iterations reproducing a certain regular pattern at finer and finer resolution. ‘Shuffles of Min’ constructions (Kimeldorf and Sampson, 1978, Mikusinski et al, 1992) are of the same nature, while Zhang (2019) constructed a similar example based on a ‘bisection expanding cross’ (Figure 5.4 below); see also Sejdinovic et al (2013, p. 2287) who consider sine curves of higher and higher frequencies. Denote by Ψ_d the approximation of the fractal Ψ at resolution $d \in \mathbb{N}$, and $C_d \subseteq I^2$ its image. For any finite d , $\iint_{C_d} dF_1(x_1)dF_2(x_2) = 0$, so that defining $(X_1, X_2) = \Psi_d(U)$ would produce a couple of purely dependent variables according to (P6). Now, in the limit $d \rightarrow \infty$, their pure dependence suddenly turns into independence, meaning that one can approach arbitrarily closely, in a certain sense, distributions showing independence by distributions showing pure dependence. This ostensible paradox is the core of the discussion in Zhang (2019).

This, however, is very similar to the following simple case of a degenerate bivariate Gaussian distribution: let $X_1 \sim N(0,1)$ and $X_2 = aX_1$, for $a \neq 0$. Then $|\rho(X_1, X_2)| = 1$ (pure dependence), including for a arbitrarily small. Yet, when $a = 0$, $\rho(X_1, X_2) = 0$ (independence). As $a \rightarrow 0$, one would thus approach

independence arbitrarily closely by pure dependence. The above described paradox of contiguity between pure dependence and independence is thus well understood.

Rényi (1959)'s original Axiom E) requires D to be maximum under 'strict dependence', defined as when there exists a Borel function $\psi_1 : \mathbb{R} \rightarrow \mathbb{R}$ such that $X_2 = \psi_1(X_1)$, or a Borel function $\psi_2 : \mathbb{R} \rightarrow \mathbb{R}$ such that that $X_1 = \psi_2(X_2)$; that is Axiom 3 in Granger et al (2004) as well. In other words, one of the variables should be deterministically predictable from the other, that is, what Lancaster (1963) defined as 'complete dependence'. Yet X_2 may be a deterministic function of X_1 while giving very little information on X_1 ; e.g., $X_2 = \sin(X_1)$ (ψ_1 is many-to-one). Clearly asymmetric, this concept seems hardly reconcilable with (P2). In an attempt to symmetrise it, one can request the existence of two functions ψ_1 and ψ_2 such that $X_2 = \psi_1(X_1)$ and $X_1 = \psi_2(X_2)$, that is, a one-to-one relationship between X_1 and X_2 . This would reduce any sense of perfect dependence to 'mutual complete dependence' (Lancaster, 1963) or even 'monotone dependence' (Kimeldorf and Sampson, 1978), which appears too restrictive. All in all, if some dependence between X_1 and X_2 may usually help for predicting X_2 from X_1 or vice-versa, the concepts of predictability and dependence are indeed distinct.

Note that the characterisation of 'pure dependence' in (P6) is symmetric in X_1 and X_2 , and it admits deterministic predictability as a particular case, in the sense that if $X_2 = \psi_1(X_1)$, one can write

$$(X_1, X_2) = (F_1^{-1}(U), \psi_1(F_1^{-1}(U))) \square \Psi(U), \quad \text{for } U \sim U_{[0,1]}.$$

Generally, though, it does not require any of the two variables to be predictable from the other. We note that many popular dependence measures fail to satisfy (P6). In particular, Pearson's correlation and distance correlation (dCor) take their maximum value only in the case of perfect linear relationship between X_1 and X_2 , while rank-based measures such as Spearman's ρ , Kendall's τ or Hoeffding's D are maximum only in the case of 'monotone dependence' (deterministic monotonic relationship).

2.4 Generalised Data Processing Inequality, equitability, margin-freeness and copulas

The *Data Processing Inequality* is an important information-theoretic concept (Cover and Thomas, 2006, Section 2.8). It carries the intuitively clear idea that information cannot be gained when a signal goes through a noisy channel. Specifically, if X_1 , X_2 and X_3 are three random variables such that X_1 and X_3 are independent conditionally on X_2 ($X_1 \perp\!\!\!\perp X_3 \mid X_2$), it establishes that $I(X_1, X_3) \leq I(X_1, X_2)$ where $I(X_i, X_j)$ is the Mutual Information between X_i and X_j ($i, j \in \{1, 2, 3\}$). Concretely, if X_2 is a signal containing some information about X_1 and we are only able to see X_3 , a version of X_2 diluted in white noise, then the noisy version X_3 is necessarily less informative about X_1 than X_2 . It seems fair to paraphrase this as ‘there is less dependence between X_1 and X_3 than between X_1 and X_2 ’, making it reasonable to ask a dependence measure to satisfy the ‘generalised Data Processing Inequality’ (P7).

The implications of (P7) are actually very deep. In particular, we have the following result:

Proposition 2.1. *A dependence measure D satisfying (P7) is such that*

$$D(X_1, X_2) = D(\psi_1(X_1), \psi_2(X_2)) \quad (2.1)$$

for any Borel functions ψ_1 and ψ_2 such that $X_1 \perp\!\!\!\perp X_2 \mid \psi_1(X_1)$ and $X_1 \perp\!\!\!\perp X_2 \mid \psi_2(X_2)$.

Proof. $X_1 \perp\!\!\!\perp X_2 \mid \psi_1(X_1) \Rightarrow D(X_1, X_2) \leq D(\psi_1(X_1), X_2)$, by (P7). Given X_1 , $\psi_1(X_1)$ is degenerated, so independent of any other random variable. Thus $\psi_1(X_1) \perp\!\!\!\perp X_2 \mid X_1$ and $D(\psi_1(X_1), X_2) \leq D(X_1, X_2)$, again by (P7). Hence $D(\psi_1(X_1), X_2) = D(X_1, X_2)$. The second part follows identically, as $X_1 \perp\!\!\!\perp X_2 \mid \psi_2(X_2) \Rightarrow \psi_1(X_1) \perp\!\!\!\perp X_2 \mid \psi_2(X_2)$. \square

This property is strongly related to the concept of equitability, which recently came to light for discriminating between dependence measures. In short, a dependence measure is equitable if it returns the same value to equally noisy relationships of different nature. After it was empirically outlined by Reshef et al (2011), several formal definitions were proposed (Kinney and

Atwal, 2014, Reshef et al, 2016, Ding et al, 2017), and (2.1) is actually a slight generalisation of Kinney and Atwal (2014)'s definition of *self-equitability*. The conditional independence assumptions $X_1 \perp\!\!\!\perp X_2 \mid \psi_1(X_1)$ and $X_1 \perp\!\!\!\perp X_2 \mid \psi_2(X_2)$ essentially mean that the whole dependence between X_1 and X_2 can be captured by the functions ψ_1 and/or ψ_2 , and the dependence measure should reflect that. This is the case, for instance, under the regression model $X_2 = \psi_1(X_1) + \varepsilon$, where $\varepsilon \perp\!\!\!\perp X_1$. Then (2.1) implies that $D(X_1, X_2) = D(\psi_1(X_1), X_2)$, meaning that the value of D is only driven by the signal-to-noise ratio, irrespective of the nature or shape of ψ_1 .

Note that the set of functions ψ_1, ψ_2 such that $X_1 \perp\!\!\!\perp X_2 \mid \psi_1(X_1)$ and $X_1 \perp\!\!\!\perp X_2 \mid \psi_2(X_2)$ is necessarily non-empty. In particular, all strictly monotonic Borel functions ψ_1 and ψ_2 are such functions, as in that case, conditioning on $\psi_1(X_1)$ or on $\psi_2(X_2)$ is equivalent to conditioning directly on X_1 or X_2 , respectively. Thus, (2.1) implies that D is invariant to monotonic transformations of X_1 and X_2 , i.e.,

$$D(\psi_1(X_1), \psi_2(X_2)) = D(X_1, X_2) \quad (2.2)$$

for any two strictly monotonic Borel functions $\psi_1, \psi_2 : \mathbb{R} \rightarrow \mathbb{R}$. This has often been presented as a fundamental trait of any valid dependence measure: it is Axiom F) in Rényi (1959)'s original paper and Axiom 6 in Granger et al (2004), for instance. A dependence measure satisfying (2.2) is said margin-free, given that the marginal distributions can be arbitrarily distorted by ψ_1 and ψ_2 without affecting its value. It so appears that (P7) is actually a more fundamental property than 'margin-freeness', given that $(P7) \Rightarrow (2.1) \Rightarrow (2.2)$.

Margin-freeness is typically associated to copulas. The copula C_{12} of the continuous vector (X_1, X_2) is the distribution of $(F_1(X_1), F_2(X_2))$ on the unit square I^2 , and is known to capture all the characteristics of F_{12} which are invariant to monotonic transformations of its margins (Schweizer and Wolff, 1981). Thus, any dependence measure which is explicitly copula-based, e.g., Spearman's ρ , Kendall's τ or Hoeffding's D , is margin-free in the continuous setting. Conversely, for X_1 and X_2 continuous, any margin-free

dependence measure $D(X_1, X_2)$ must admit a representation involving only the copula of (X_1, X_2) . Many popular measures of dependence are not copula-based hence they are not margin-free. Besides the obvious example of Pearson's correlation, it is the case of the distance correlation dCor ([Székely et al, 2007](#)) and the Hilbert-Schmidt independence criterion HSIC ([Gretton et al, 2005](#)), among others. Problems that this creates were implicitly acknowledged by [Székely and Rizzo \(2009, Section 4.3\)](#) as they briefly mentioned basing empirical estimation of dCor on the ranks of the observations instead of on the original observations. Likewise, [Póczos et al \(2012\)](#) suggested a copula version of HSIC. Here it is stressed that, beyond margin-freeness, the real issue with measures not copula-based is that they violate (P7), whose legitimacy seems difficult to contest.

3 ϕ -dependence measures

3.1 Definition and properties

A measure of the dependence in (X_1, X_2) should quantify how much different is the joint distribution F_{12} from the product $F_1 F_2$ of its marginals; see Section 2.1. Natural candidates for this task are the ϕ -divergences ([Ali and Silvey, 1966](#)) between F_{12} and $F_1 F_2$, viz.

$$\Delta_\phi(F_{12} \parallel F_1 F_2) = \iint_{\mathbb{R}^2} \phi\left(\frac{dF_{12}(x_1, x_2)}{dF_1(x_1)dF_2(x_2)}\right) dF_1(x_1)dF_2(x_2) \quad (3.1)$$

for some convex function ϕ such that $\phi(1) = 0$. The ϕ -divergence family includes most of the common statistical distances between distributions ([Liese and Vajda, 2006](#)), hence (3.1) includes many familiar dependence measures. For instance, $\phi(t) = t \log t$ yields the Kullback-Leibler divergence between F_{12} and $F_1 F_2$, which is the Mutual Information $I(X_1, X_2)$.

Provided that we allow the Radon-Nikodym derivative $dF_{12} / dF_1 dF_2$ to be infinite in case of singularity (see discussion in [Silvey \(1964\)](#) and Remark 3.1 below), $\Delta_\phi(F_{12} \parallel F_1 F_2)$ is always defined and 'existence' (P1) is guaranteed. 'Symmetry' (P2) is obvious from the definition. 'Weak Gaussian conformity' (P5) is what [Ali and Silvey \(1965\)](#) established. 'Generalised Data Processing

Inequality' (P7) holds by Theorem 4 of [Kinney and Atwal \(2014\)](#), which makes any measure (3.1) automatically margin-free. Indeed, with X_1 and X_2 both continuous, one has the copula form

$$\Delta_\phi(F_{12} \parallel F_1 F_2) = \iint_{I^2} \phi\left(\frac{dC_{12}(u_1, u_2)}{du_1 du_2}\right) du_1 du_2 = \iint_{I^2} \phi(c_{12}(u_1, u_2)) du_1 du_2, \quad (3.2)$$

through Probability Integral Transform $u_k = F_k(x_k)$, $k = 1, 2$, where C_{12} is the copula of (X_1, X_2) and c_{12} its density.

Remark 3.1. The copula density c_{12} is obviously defined when C_{12} is absolutely continuous. If C_{12} is singular or has a singular component, then one can think of c_{12} as infinite on the singularity, and defined as the limit of the densities of a sequence of absolutely continuous copulas converging to C_{12} in an appropriate sense ([Ding et al, 2017](#), p. 9).

There remain (P3), (P4) and (P6), compliance to which depends on ϕ as follows.

Proposition 3.1. *Let $\phi: (0, +\infty) \rightarrow \mathbb{R}$ be convex with $\phi(1) = 0$, and call $\varphi_0 = \lim_{t \rightarrow 0} \phi(t)$, $\varphi_0^* = \lim_{t \rightarrow 0} t\phi(1/t)$ and $\bar{\varphi} = \varphi_0 + \varphi_0^*$. Then, $\Delta_\phi(F_{12} \parallel F_1 F_2)$ in (3.1) is such that $0 \leq \Delta_\phi(F_{12} \parallel F_1 F_2) \leq \bar{\varphi}$. In addition, if ϕ is strictly convex at $t = 1$, then $\Delta_\phi(F_{12} \parallel F_1 F_2) = 0$ if and only if $X_1 \perp\!\!\!\perp X_2$ and, provided that $\bar{\varphi} < \infty$, $\Delta_\phi(F_{12} \parallel F_1 F_2) = \bar{\varphi}$ if and only if X_1 and X_2 are purely dependent in the sense of (P6).*

Proof. This follows from Theorem 5 in [Liese and Vajda \(2006\)](#). \square

'Characterisation of independence' (P4) is thus granted as soon as ϕ is strictly convex at $t = 1$. If $\bar{\varphi} < \infty$, then 'normalisation' (P3) is achieved by obvious linear rescaling $\Delta_\phi(F_{12} \parallel F_1 F_2) \square (1/\bar{\varphi})\Delta_\phi(F_{12} \parallel F_1 F_2)$, and 'characterisation of pure dependence' (P6) for $\Delta_\phi(F_{12} \parallel F_1 F_2)$ follows straight from Proposition 3.1. On the other hand, if $\bar{\varphi} = \infty$, then, even though one can still enforce (P3) through a non-linear transformation, (P6) cannot be made true: there are cases in which $\Delta_\phi(F_{12} \parallel F_1 F_2)$ is infinite while X_1 and X_2 do not

show any sense of strong dependence (see Section 3.2). In short, a ϕ -dependence measure fails to comply with (P6) when its baseline version (3.1) can be infinite – see comments in [Micheas and Zografos \(2006\)](#) (in particular, their Axiom 9) – which has also serious implications when empirically estimating such a measure ([Ding et al, 2017](#), Section 3.1).

3.2 Common choices

If one takes $\phi(t) = (t-1)^2$, for which choice $\bar{\varphi} = \infty$, $\Delta_\phi(F_{12} \| F_1 F_2)$ is Pearson's Mean Square Contingency coefficient $\Phi^2(X_1, X_2)$, allowed to be infinite. It is usually re-normalised as $\Phi^{\hat{}}(X_1, X_2) = \Phi(X_1, X_2) / (1 + \Phi^2(X_1, X_2))^{1/2} \in [0, 1]$ so as to agree with $|\rho|$ in the Gaussian case ([Rényi, 1959](#)). However, in general this may equal 1 even when X_1 and X_2 do not show any sense of strong relationship. Indeed, in the form (3.2), $\Phi^2(X_1, X_2) = \iint_{I^2} (c_{12}(u_1, u_2) - 1)^2 du_1 du_2$, and it is known that the copula density c_{12} is not square-integrable as soon as there is the *slightest* level of tail dependence between X_1 and X_2 ([Beare, 2010](#), Theorem 3.3). This observation was already used by [Hoeffding \(1942, p. 150\)](#) for calling into question the usefulness of Φ^2 as a measure of dependence.

Likewise, for $\phi(t) = t \log t$, one gets $\bar{\varphi} = \infty$, and indeed the Mutual Information $I(X_1, X_2)$ can be infinite. It can be re-normalised as $I^{\hat{}}(X_1, X_2) = [1 - \exp(-2I(X_1, X_2))]^{1/2} \in [0, 1]$ so as to agree again with $|\rho|$ in the Gaussian case ([Linfoot, 1957](#)). Yet, it cannot unequivocally characterise pure dependence. For instance, for some $a \in (0, 1)$ and $m \in \mathbb{N}$, consider the copula density

$$c(u, v) = \begin{cases} \frac{1}{1-a} & \text{if } (u, v) \in [0, 1-a]^2 \\ \frac{m}{a} & \text{if } (u, v) \in [1-a + (v-1)a/m, 1-a + va/m]^2, v=1, 2, \dots, m, \\ 0 & \text{elsewhere.} \end{cases}$$

Then it can be checked that

$$I(X_1, X_2) = \iint_{\mathbb{I}^2} c(u_1, u_2) \log c(u_1, u_2) du_1 du_2 = -(1-a) \log(1-a) - a \log a + a \log m.$$

As $a \rightarrow 0$, the copula density is constant over a set arbitrarily close to the whole \mathbb{I}^2 , suggesting a situation of near-independence. However, $I(X_1, X_2)$ can be made arbitrarily large by growing m such that $a \log m \rightarrow \infty$.

By contrast, the choice $\phi(t) = |t-1|$ yields $\bar{\phi} = 2 < \infty$. The associated (rescaled) dependence measure $\Delta_\phi^\Delta(F_{12} \parallel F_1 F_2) = (1/2) \Delta_\phi(F_{12} \parallel F_1 F_2)$ consequently satisfies (P1)–(P7). It is actually Silvey (1964)'s Δ , renamed 'robust copula dependence' measure in Ding et al (2017) when in copula form (3.2).

However, no root- n consistent estimator of that measure seems available. In the next section, we explore a particular measure which lies at the intersection between two common families of ϕ -divergences (see Appendix A), and we provide a simple root- n consistent estimator.

4 The Hellinger correlation

If one takes $\phi(t) = (t^{1/2} - 1)^2$, for which $\bar{\phi} = 2 < \infty$, the corresponding (rescaled) measure

$$\begin{aligned} \Delta_\phi^\Delta(F_{12} \parallel F_1 F_2) &= \frac{1}{2} \iint_{\mathbb{I}^2} \left(\sqrt{\frac{dF_{12}(x_1, x_2)}{dF_1(x_1) dF_2(x_2)}} - 1 \right)^2 dF_1(x_1) dF_2(x_2) \\ &= \frac{1}{2} \iint_{\mathbb{I}^2} \left(\sqrt{dF_{12}(x_1, x_2)} - \sqrt{dF_1(x_1) dF_2(x_2)} \right)^2 \square H^2(X_1, X_2) \end{aligned} \quad (4.1)$$

satisfies (P1)–(P7). We denote this measure $H^2(X_1, X_2)$, or simply H^2 , as it is the squared Hellinger distance between F_{12} and $F_1 F_2$. Under the copula form (3.2), it is

$$H^2 = \frac{1}{2} \iint_{\mathbb{I}^2} \left(\sqrt{c_{12}}(u_1, u_2) - 1 \right)^2 du_1 du_2 = 1 - \iint_{\mathbb{I}^2} \sqrt{c_{12}}(u_1, u_2) du_1 du_2 \square 1 - B, \quad (4.2)$$

where B is the *Bhattacharyya affinity coefficient* (Bhattacharyya, 1943) between the copula density c_{12} and the independence copula density identically equal to 1 on \mathbb{I}^2 .

In the bivariate Gaussian case, it can be checked that

$H^2 = 1 - (2(1 - \rho^2)^{1/4}) / (4 - \rho^2)^{1/2} \square h(\rho)$, a strictly increasing function of the correlation ρ – in agreement with ‘weak Gaussian conformity’ (P5). This suggests to consider the measure on the transformed scale $\eta = h^{-1}(H^2)$. As h^{-1} is a bijection from $[0,1]$ to $[0,1]$, it preserves all (P1)–(P7) for η . Direct algebra yields

$$\eta = (2/B^2) \{B^4 + (4 - 3B^4)^{1/2} - 2\}^{1/2}, \quad (4.3)$$

our proposed measure of dependence. We call $\eta \square \eta(X_1, X_2)$ the *Hellinger correlation* between X_1 and X_2 , given that it is defined so that $\eta = |\rho|$ in the bivariate Gaussian case (‘strong Gaussian conformity’). This greatly facilitates interpretation as the value of η can easily be appreciated on the familiar Pearson’s correlation scale: a Hellinger correlation η equal to $\eta_0 \in [0,1]$ represents a dependence of the same strength as in a bivariate Gaussian vector whose Pearson’s correlation is $\rho = \eta_0$.

5 Empirical estimation

5.1 Background

Measuring dependence by H^2 has been considered before, e.g., by Granger et al (2004) who proposed an estimator based on kernel density estimation and numerical integration – see also Su and White (2008). However, the obtained estimate heavily depends on the bandwidths used in the kernel estimators (Skaug and Tjøstheim, 1996), an appropriate choice of which in practice remaining problematic. Rosenblatt (1975), Hong and White (2005) and Ding et al (2017) face the same issue when estimating their proposed measure; see also comments in Zeng et al (2018). Of course, difficulty in estimating a measure from empirical data seriously limits its practical reach.

Now, if knowledge of dF_{12} , dF_1 and dF_2 in (4.1), or of the copula density c_{12} in (4.2), implies knowledge of H^2 , the contrary is not true: one cannot recover dF_{12} , dF_1 and dF_2 or c_{12} from H^2 alone. Empirical estimation of H^2 (or any function thereof, such as B or η) should thus not be based on the more difficult task of estimating the individual densities. Indeed it has been known,

at least since [Kozachenko and Leonenko \(1987\)](#), that one can consistently estimate some ϕ -divergences without consistently estimating the underlying distributions; [Berrett et al \(2018\)](#) offer a recent review of such ideas for entropy estimation. Below we suggest a simple estimator of $B = \iint \sqrt{c_{12}}$ along that line, subsequently producing an estimator of η through (4.3).

5.2 Basic estimator and asymptotic properties

Let $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}$, where $\mathbf{X}_i = (X_{i1}, X_{i2}) \in \mathbb{I}^2$, be a random sample from F_{12} . For $i \in \{1, \dots, n\}$, call $\mathbf{U}_i = (U_{i1}, U_{i2}) = (F_1(X_{i1}), F_2(X_{i2}))$, the corresponding observations from the copula C_{12} , and $\mathbf{U}_i = (\hat{U}_{i1}, \hat{U}_{i2}) = (\hat{F}_{n,1}(X_{i1}), \hat{F}_{n,2}(X_{i2}))$ the ‘pseudo’-observations obtained from $\hat{F}_{n,k}(x) = (1/(n+1)) \sum_{i=1}^n I_{\{X_{ik} \leq x\}}$, $k = 1, 2$, as is customary in the copula literature. Let $R_i = \min_{j \neq i} \|\mathbf{U}_j - \mathbf{U}_i\|_2$, the Euclidean distance between \mathbf{U}_i and its closest neighbour, and its ‘pseudo’-version $\hat{R}_i = \min_{j \neq i} \|\mathbf{U}_j - \mathbf{U}_i\|_2$.

Our first simple and entirely data-based estimator for B , not involving any user-defined parameter, is

$$B = \frac{2\sqrt{n-1}}{n} \sum_{i=1}^n \hat{R}_i, \quad (5.1)$$

which is the feasible version of the oracle estimator $B \approx (2\sqrt{n-1}/n) \sum_{i=1}^n R_i$. See

[Leonenko et al \(2008, Remark 3.1\)](#) for the motivation behind this estimator. In terms of measuring dependence, the intuition is the following. In case of pure dependence ($B = 0$), the \mathbf{U}_i ’s fall exactly on a curve. The R_i ’s are then essentially univariate spacings, known to be of order $O_p(n^{-1})$ ([Pyke, 1965](#)).

This yields $B \rightarrow 0$ in probability. Now gradually relaxing dependence amounts to allowing some play around that curve for the \mathbf{U}_i ’s. As these get more and more room to move apart, the R_i ’s globally increase, and so does B .

Ultimately, when the \mathbf{U}_i ’s get totally free (independence, $B = 1$), they uniformly cover \mathbb{I}^2 and maximally occupy their allowed space. The R_i ’s are globally as large as can be, and so is B . The latitude given to the \mathbf{U}_i ’s for

covering I^2 reflects the number of degrees of freedom in (X_1, X_2) ; see Section 2.3.

The L_2 -consistency of B follows straight from [Leonenko et al \(2008, Section 3.1\)](#). [Aya Moreno et al \(2018, Lemma 1\)](#) and [Ebner et al \(2018, Theorems 1 and 2\)](#) further establish the root- n consistency and the asymptotic normality of B , if the copula density c_{12} is bounded on I^2 . In Appendix B we show how to relax this restrictive assumption by combining marginal transformations ([Geenens et al, 2017](#)) and results of [Singh and Póczos \(2016\)](#). [Deheuvels \(2009, Lemmas 2.1 and 2.2\)](#) help to establish that $B - \hat{B} = O_p(n^{-1/2})$, meaning that the above root- n consistency of the oracle estimator B carries over to the feasible \hat{B} .

5.3 Normalisation

In finite samples, however, \hat{B} suffers from two serious defects (as would B). First, it is heavily biased due to the boundedness of the support I^2 of C_{12} ([Liitiäinen et al, 2010](#)). Second, although $B = \iint \sqrt{c_{12}} \leq \iint (\sqrt{c_{12}})^2 = 1$ by Cauchy-Schwarz, it may happen that $\hat{B} > 1$, implying a meaningless negative estimate of H^2 and precluding ulterior estimation of η .

Now, evidently $\sqrt{c_{12}} \in L_2(I^2)$. Hence from any orthonormal basis of $L_2(I)$, say $\{b_0(u), b_1(u), \dots\}$ with $b_0(u) \equiv 1$, one can form a tensorised orthonormal basis for $L_2(I^2)$ and write the expansion $\sqrt{c_{12}}(u_1, u_2) = \sum_{k=0}^{\infty} \sum_{\ell=0}^{\infty} \beta_{k\ell} b_k(u_1) b_{\ell}(u_2)$, where $\beta_{k\ell} = \iint_{I^2} \sqrt{c_{12}}(u_1, u_2) b_k(u_1) b_{\ell}(u_2) du_1 du_2$. In particular, see that $B = \beta_{00}$. Similarly to (5.1),

$$\hat{\beta}_{k\ell} = \frac{2\sqrt{n-1}}{n} \sum_{i=1}^n \hat{R}_i b_k(\hat{U}_{i1}) b_{\ell}(\hat{U}_{i2}), \quad k, \ell \in \mathbb{N}, \quad (5.2)$$

is a root- n consistent estimator of $\beta_{k\ell}$ ([Aya Moreno et al, 2018, Lemma 1](#)). For appropriate cut-off values K and L , define $\hat{\sqrt{c}}_{12}(u_1, u_2) = \sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell} b_k(u_1) b_{\ell}(u_2)$ and

orthogonal series estimator for $\sqrt{c_{12}}$, and see that

$$\iint_{I^2} \left(\sqrt{c_{12}} \right)^2 (u_1, u_2) du_1 du_2 = \sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^2. \text{ This suggests to re-normalise } B = \hat{\beta}_{00} \text{ as}$$

$$B_{KL} = B / \left(\sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^2 \right)^{1/2}. \quad (5.3)$$

Not only this shrinkage guarantees $B_{KL} \in [0,1]$ always, it reduces the variance and mostly takes care of the boundary effect as well. Intuitively, when U_i is close to the boundary of I^2 , its neighbourhood is empty of data on one side and \hat{R}_i tends to be larger than what it should be. This makes (5.1) overestimate B and induces bias. But then (5.2) overestimates each $\beta_{k\ell}$ to the same extent, and the induced biases mostly cancel each other out in the ratio (5.3). Finally, (5.3) is plugged into (4.3) to define the estimator

$$\hat{\eta}_{KL} = (2 / B_{KL}^2) \left\{ B_{KL}^4 + \left(4 - 3B_{KL}^4 \right)^{1/2} - 2 \right\}^{1/2}, \quad K, L \in \mathbb{N}. \quad (5.4)$$

The following simulation showcases the performance of this estimator.

Random samples of size $n = 500$ were generated from a bivariate Gaussian distribution with correlation $\rho = 0.4$ and $\rho = 0.8$. On each of them, the basic estimator $\hat{\eta}_0$, obtained by plugging B in (4.3), was computed, together with normalised versions $\hat{\eta}_{11}$, $\hat{\eta}_{22}$ and $\hat{\eta}_{33}$. The basis $\{b_0(u), b_1(u), \dots\}$ was formed by the normed Legendre polynomials shifted to $[0,1]$. The returned estimates are shown by the boxplots in Figure 5.1. In the case $\rho = 0.4$, though, around 56% of the initially returned estimates B were found greater than 1, precluding estimation of η . So the boxplot at the extreme left only represents the 44% of the estimates which could be computed.

The reduction in both bias and variance brought by the normalisation is obvious. For $\rho = 0.4$, one extra term in each direction ($K = L = 1$) in the expansion for $\sqrt{c_{12}}$ is enough, as $\sqrt{c_{12}}$ is rather flat and well approximated by a low-degree polynomial. For $\rho = 0.8$, slightly more terms should be included as $\sqrt{c_{12}}$ tends to peak in the corners of I^2 . In order to keep the estimator totally data-driven, we suggest in Appendix C a novel and explicit cross-

validation criterion easy to minimise. The estimator computed with the values of K and L minimising that criterion is marked $\hat{\eta}_{cv}$ in Figure 5.1. The proposed cross-validation criterion consistently identifies the suitable level of approximation K and L and produces a reliable estimator of η . Estimators (5.2)-(5.3)-(5.4), as well as the suggested cross-validation procedure, have been efficiently implemented in the freely available R package `HellCor`.

For comparison, the empirical Pearson's correlation $\hat{\rho}$ was computed on the same samples, being here the 'gold standard' given that in the considered bivariate Gaussian vectors, $\eta = \rho$. Interestingly, it is seen that $\hat{\eta}_{cv}$ is only marginally dominated by $\hat{\rho}$, the classical estimator specifically designed for capturing the dependence of linear nature peculiar to bivariate Gaussian vectors.

Naturally, the Hellinger correlation would capture dependence of any other nature as well. This is illustrated in Figure 5.2, showing 15 random samples of size $n = 500$ generated from the 15 scenarii considered in [Heller et al \(2016, Figure 4\)](#). The estimator $\hat{\eta}_{cv}$ with expansion in the Legendre basis and cross-validation cut-offs, from now on denoted simply $\hat{\eta}$, was computed on each of them. The estimated values of η are shown on top of each plot, attesting an interesting descriptive ability. In Appendix D, these 15 examples, placed in a much broader framework, are the basis of a comprehensive study of the accuracy of our proposed estimator $\hat{\eta}$ of η and the power of an independence testing procedure based on it.

Finally, we generated random samples of size $n = 500$ and $n = 5,000$ on the Peano curve (Figure 5.3) and on [Zhang \(2019\)](#)'s 'bisection expanding cross' (Figure 5.4), at increasing resolution d (see Section 2.3). For any finite d , the joint distribution of (X_1, X_2) is concentrated on a curve (in the sense of (P6)) and hence exhibits perfect dependence (population Hellinger correlation $\eta = 1$). However, that dependence is harder and harder to detect empirically as the resolution increases: for the same sample size, the length of the 'curve' increases, hence the points are more distant to one another and it becomes difficult to capture their perfect alignment. This explains why the empirical

Hellinger correlation tends to fade as d increases (for n fixed). Increasing the sample size pushes the empirical Hellinger correlation towards $\eta = 1$ for any finite d , and towards $\eta = 0$ for $d = \infty$ – as explained in Section 2.3, X_1 and X_2 are indeed independent at the limit $d = \infty$. For comparison, we have also computed the empirical distance correlation ($dCor$) for each of those samples. The observed values of $dCor$ remain very low across all scenarii, actually close to its value for the independent sample ($d = \infty$), even for $n = 5,000$. The observed value is significantly different to 0 (at level $\alpha = 0.05$) only for resolution $d = 1$, and $d = 2$ for $n = 5,000$ in the ‘bisection expanding cross’ case (obtained from 200 permutations, from `dcor.test` in the R package `energy`). We also computed a variety of other dependence measures or dependence test statistics (not shown), such as Hoeffding’s D (Hoeffding, 1948), HSIC (Pfister et al, 2018), or the HHG test statistic (Heller et al, 2016), with essentially the same results: none is able to detect dependence beyond resolution $d = 1$. The power of $\hat{\eta}$ to do so is showcased in Section 5.5.

5.4 Accuracy of empirical estimation

Figure 5.2 illustrates the descriptive ability of $\hat{\eta}$ on a given sample, however it does not allow us to judge the accuracy of that empirical estimator for the population value η . Indeed, in those 15 scenarii, the underlying copula (hence its density) is not known, and the ‘true’ value η is not available. So, to investigate how close our proposed estimator $\hat{\eta}$ is to η , we consider here nine common one-parameter copula families (see Figure 5.5) whose distributions are available in closed form: the respective true values of B can be computed by direct numerical integration in (4.2), and then plugged into (4.3) to obtain the true η ’s. (For this task we used the R package `copula`, and we double-checked the results with `Mathematica`. We requested a relative accuracy of 10^{-4} for η .) Figure 5.5 shows the value of η as a function of the parameter θ classically parameterising each of those copula families. Naturally, the function $\eta(\theta)$ for the Gaussian copula is the absolute value function, as the Hellinger correlation η has been calibrated for this to be the case (θ is Pearson’s correlation for the Gaussian copula). Other notable observations

can be made from Figure 5.5, e.g., the fact that the Student- t copula can never describe independence.

For the below simulation, we chose one particular value of θ for each family, so that, overall, we have a wide range of associated values for η , from 'low' to 'high'; see Table 5.1.

For each of these copulas, we generated $M = 1,000$ samples of sizes $n = 100, 250, 500$ and $1,000$ from which we computed $\hat{\eta}$. The associated boxplots are shown in Figure 5.6, empirically confirming the aptitude of the proposed estimator $\hat{\eta}$ for accurately approximating η as the sample size increases. A slight positive bias for small values of n is also apparent for the cases considered here. Further results of this type are provided in Appendix D.

5.5 Significance and independence testing

The *exact* statistical significance of the empirical Hellinger correlation computed on a given data set can easily be obtained for any sample size by Monte-Carlo simulations. Indeed, if X_1 and X_2 are independent, then C_{12} is bivariate uniform on I^2 . One can then simulate a large number of bivariate uniform samples of any size n and compute $\hat{\eta}$ on each of those, which would allow arbitrary close approximation to the exact sampling distribution of $\hat{\eta}$ for that n in the case of independent variables ($\eta = 0$). For instance, from $M = 10,000$ independent bivariate uniform samples of size $n = 500$ and $n = 5,000$, we have obtained, respectively, critical values $\bar{\eta}_c = 0.146$ and $\bar{\eta}_c = 0.048$ (exact up to Monte Carlo error) at significance level $\alpha = 0.05$. All the observed empirical Hellinger correlations in Figure 5.2 are thus statistically significant, except under the '4 clouds' scenario [6] in which case X_1 and X_2 are indeed independent (Heller et al, 2016). For the Peano and 'bisection expanding cross' scenarii (Figures 5.3 and 5.4), dependence is detected up to resolution $d = 3$ for $n = 500$. At that sample size, it becomes hardly possible to visually make any distinction between the samples for $d = 4, 5$ and $d = \infty$ (independence), indeed. For $n = 5,000$, on the other hand, dependence is

detected up to resolution $d = 5$. This can be contrasted to the observations made at the end of Section 5.3 about other tests or measures.

Naturally, checking the significance (i.e., non-nullity) of a measure of dependence is in many ways akin to testing for independence. Now, bearing in mind the discussion in Section 1, it seems clear that $\hat{\eta}$ might not be a test statistic leading to the most powerful test. Yet, the previous results indicate that, in some situations, the Hellinger correlation may still show some interesting ability to highlight dependence compared to other, more classical choices of dependence measure. We investigate this further here.

We generated $M = 10,000$ samples of increasing sizes for the Peano ($d = 4$) and ‘bisection expanding cross’ ($d = 3$) structures of dependence, respectively (see Figures 5.3 and 5.4). From these samples, we computed the (empirical) power of independence tests using for test statistics: (i) $\hat{\eta}_0$, the ‘raw’ empirical Hellinger correlation, without any normalisation; and (ii) $\hat{\eta}$, its properly normalised version (with expansion in the Legendre basis and cross-validation cut-offs), as described in Section 5.3. Exact critical values (as opposed to asymptotic or based on permutations) for these two tests were obtained for all n as described above. We compared these powers to those of five popular competitors, namely dCov: distance correlation ([Székely et al, 2007](#)); HHG: [Heller et al \(2012\)](#); G^2 : [Wang et al \(2017\)](#); RCD: robust copula dependence ([Ding et al, 2017](#)); GMR: [Granger et al \(2004\)](#). The results are shown in Tables 5.2 and 5.3.

For the Peano scenario, all competitors stay stuck at the marginal power ($\sim 5\%$), even for ‘large’ samples ($n = 500$). For the ‘bisection expanding cross’ scenario, dCov and G^2 fail to detect any dependence again, while RCD, GMR and to a lesser extent HHG manage to do so, although with a power much lower than the test based on $\hat{\eta}_0$. Note that the resolution $d = 4$ for Peano being higher than for ‘bisection expanding cross’ ($d = 3$), the dependence in this second scenario is more obvious than in the first one, which explains higher power altogether across the table.

What appears clearly is that the test based on $\hat{\eta}_0$ (non-normalised version of the Hellinger correlation) is the most powerful here, and by far. In particular, it has a much higher power than the test based on the normalised one $\hat{\eta}$ (which is still more powerful overall than most of the competitors). This illustrates our point, discussed in Section 1, that a good measure of dependence does not necessarily produce a good test, at least not the best test: the normalisation was introduced in Section 5.3 for stabilising the variance of the ‘raw’ estimator $\hat{\eta}_0$ and forcing it to belong to $[0,1]$, hence for producing a much more meaningful measure – but this apparently comes at the price of a lower power when it comes to test for independence.

The above scenarii are two of ‘pure dependence’ – although their particular structure makes it very difficult to detect. In Appendix D we provide a more comprehensive power study where different levels of dependence are considered. Indeed the tests based on $\hat{\eta}_0$ and $\hat{\eta}$ are not always the most powerful, but they consistently rank in the best options overall.

6 Real data applications

6.1 Coral fish, seabirds and reef productivity

We consider data from a recent study on coral reef productivity described in [Graham et al \(2018\)](#). The density of some species of seabirds and coral reef fish around $n = 12$ islands of the Chagos Archipelago (British Indian Ocean Territory) was recorded; see Table 6.1. Those seabirds forage and feed in the open ocean, far from reefs, and their number around a given island largely depends on the presence or absence of rats on it. So they belong to a different ecosystem to the fish. Indeed, empirical Pearson’s correlation between the densities of seabirds and fish around the 12 islands under study is $\hat{\rho} = 0.374$, not significantly different to 0 ($p = 0.231$). Likewise, the empirical distance correlation ([Székely et al, 2007](#)) is $dCor = 0.526$, not significant ($p = 0.159$ based on 200 permutations, from `dcor.test` in the R package `energy`).

This failure to evidence any significant dependence goes, however, against the report of [Graham et al \(2018\)](#), whose main finding is that the two

ecosystems are in fact connected. Nutrients, in particular nitrogen, leach from seabird guano onshore to nearshore marine systems through rainfall, among others. With this extra nutrient supply, benthic algae develop more on coral reefs adjacent to islands where seabirds are abundant, making the reef-fish communities there more abundant as well, given that fish mostly feed on those algae.

The nitrogen input, as described in [Graham et al \(2018\)](#), is thus a latent variable, positively associated to both seabird and fish densities. Now, by design (see Section 2.3), the Hellinger correlation should be particularly effective at highlighting dependence when induced by such a hidden effect. Indeed the empirical Hellinger correlation (5.4), with Legendre basis and $K = 1$ and $L = 1$ determined by cross-validation in (5.3), is here $\hat{\eta} = 0.744$. The exact p -value of significance, obtained from $M = 10,000$ independent bivariate uniform samples of size $n = 12$, as described in Section 5.5, is found to be 0.045. So, at the typical significance level $\alpha = 0.05$, the Hellinger correlation is able to highlight the dependence between ‘seabirds’ and ‘fish’, even from such a small sample, corroborating [Graham et al \(2018\)](#)’s findings.

One can also build a bootstrap confidence interval for η by resampling from the empirical beta copula ([Segers et al, 2017](#)). The conventional bootstrap (sampling the pairs with replacement from $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}$) is not appropriate here as the bootstrap resamples would typically include several times the same \mathbf{X}_i . The associated \hat{R}_i would then be 0, which would lead to gross underestimation of B in (5.1). Indeed it is known that the conventional bootstrap is inconsistent in case of estimators being non-smooth functional of F_{12} , like estimators involving nearest-neighbour distances. On the other hand, sampling from the empirical beta copula can be regarded as a variant of the ‘smoothed’ bootstrap ([Kiriliouk et al, 2019](#)). The empirical beta copula being continuous, the bootstrap resamples do not contain any repeated observation, leading to a consistent procedure. The variance of estimator (5.1) is not easily tractable, though ([Ebner et al, 2018](#), p. 235). Hence we opt for the double bootstrap procedure described in [Karlsson and Löthgren \(2000\)](#). The

produced two-sided bootstrap- t confidence intervals are known to be second-order accurate (Hall, 1992, Section 3.5). For the current data set, the obtained 95%-confidence interval for η is $[0.547, 1]$.

6.2 World demographics

As a second example we study the dependence between $X_1 = \text{'death rate'}$ and $X_2 = \text{'birth rate'}$ in 229 countries and territories around the world in the first trimester of 2020.¹ These two variables are the number of births and deaths per year per 1,000 individuals in the country. We may, indeed, expect a rather strong association between these two variables; however Figure 6.1 tells that this relationship is complex, with the cloud of points showing a clear C-shape. In fact, there seems to be two opposite effects. First, an 'decreasing' effect, from 'moderate' birth rate to 'low' birth rate as death rate increases, for the industrialised countries (mostly North America, Europe, East Asia, Australia); and second, an 'increasing' effect, from 'moderate' birth rate to 'high' birth rate as death rate increases, mostly for the African countries – this second effect being more diffused than the first one, though.

The strength of this non-linear structure of dependence is hardly captured by Pearson's correlation ($\hat{\rho} = -0.13$, not significant at level $\alpha = 0.05$). On the other hand, we get $\hat{\eta} = 0.69$ ($p = 0$), with 95% confidence interval $[0.474, 0.746]$ obtained by the bootstrap procedure described in the previous section. This illustrates the benefit of using a meaningful measure of dependence able to capture non linear relationships.

7 Perspectives

The Hellinger correlation coefficient defined in this paper only applies to the case of univariate marginal distributions. A natural generalisation is to consider the case of multivariate marginals. Specifically, let $\mathbf{X}_1 \in \mathbb{R}^{d_1}$ and $\mathbf{X}_2 \in \mathbb{R}^{d_2}$, $d_1, d_2 > 1$, with continuous distributions F_1 and F_2 respectively. Let $\mathbf{U}_1 = F_1(\mathbf{X}_1)$ and $\mathbf{U}_2 = F_2(\mathbf{X}_2)$ (componentwise), their copula transforms. Let c_{12} be the $d_1 + d_2$ -dimensional copula density of $(\mathbf{X}_1, \mathbf{X}_2)$ on $\mathcal{I}^{d_1+d_2}$ (joint density of $(\mathbf{U}_1, \mathbf{U}_2)$), c_1 be the d_1 -dimensional copula density of \mathbf{X}_1 on \mathcal{I}^{d_1}

(joint density of \mathbf{U}_1) and c_2 be the d_2 -dimensional copula density of \mathbf{X}_2 on I^{d_2} (joint density of \mathbf{U}_2). Then a marginal-free measure of dependence between the vectors \mathbf{X}_1 and \mathbf{X}_2 is the Hellinger distance between the copula C_{12} and the product of the marginal copulas $C_1 C_2$, i.e.,

$H^2 = \frac{1}{2} \iint_{I^{d_1+d_2}} \left(\sqrt{c_{12}}(\mathbf{u}_1, \mathbf{u}_2) - \sqrt{c_1}(\mathbf{u}_1) \sqrt{c_2}(\mathbf{u}_2) \right)^2 d\mathbf{u}_1 d\mathbf{u}_2$. By Parseval identity, this is

$$H^2 = \frac{1}{(2\pi)^{d_1+d_2}} \iint_{\square^{d_1+d_2}} |\varphi_{12}(\mathbf{s}, \mathbf{t}) - \varphi_1(\mathbf{s}) \varphi_2(\mathbf{t})|^2 d\mathbf{s} d\mathbf{t} \quad (7.1)$$

where for $\mathbf{s} \in \square^{d_1}$ and $\mathbf{t} \in \square^{d_2}$ and

$$i = \sqrt{-1}, \varphi_{12}(\mathbf{s}, \mathbf{t}) = \iint_{I^{d_1+d_2}} e^{i(\mathbf{s}'\mathbf{u}_1 + \mathbf{t}'\mathbf{u}_2)} \sqrt{c_{12}}(\mathbf{u}_1, \mathbf{u}_2) d\mathbf{u}_1 d\mathbf{u}_2, \varphi_1(\mathbf{s}) = \iint_{I^{d_1}} e^{i\mathbf{s}'\mathbf{u}_1} \sqrt{c_1}(\mathbf{u}_1) d\mathbf{u}_1$$

and $\varphi_2(\mathbf{t}) = \iint_{I^{d_2}} e^{i\mathbf{t}'\mathbf{u}_2} \sqrt{c_2}(\mathbf{u}_2) d\mathbf{u}_2$, that is, the Fourier transforms of $\sqrt{c_{12}}$, $\sqrt{c_1}$ and $\sqrt{c_2}$, respectively. Expanding the modulus in (7.1), one obtains

$$H^2 = 1 - \frac{1}{(2\pi)^{d_1+d_2}} \Re \left\{ \iint_{\square^{d_1+d_2}} \varphi_{12}(\mathbf{s}, \mathbf{t}) \varphi_1^*(\mathbf{s}) \varphi_2^*(\mathbf{t}) d\mathbf{s} d\mathbf{t} \right\}, \quad (7.2)$$

where \cdot^* denotes complex conjugation and $\Re(\cdot)$ is the real part. The three Fourier transforms φ_{12} , φ_1 and φ_2 are integrals with respect to the square-root of a copula density. As such, they can be estimated in a way very similar to estimating B by (5.1) or $\beta_{k\ell}$ by (5.2) in the 2-dimensional case, making use of nearest-neighbour distances in the copula-transformed domain. Lemma 1 in [Aya Moreno et al \(2018\)](#) and Theorems 1 and 2 in [Ebner et al \(2018\)](#) would again guarantee the consistency of those estimators which, when plugged in (7.2), would produce a consistent estimator of H^2 for the multivariate-marginal case. This will be investigated in detail in a follow-up paper.

Another obvious question is how to adapt the Hellinger correlation to the case of discrete variables. In fact, (4.1) obviously applies to the discrete case as well. If (X_1, X_2) is a discrete random vector with joint probability mass function p_{12} on $S_1 \times S_2$, where S_1 and S_2 are two discrete sets, and marginal distributions p_1 and p_2 , respectively, then

$$H^2(X_1, X_2) = \frac{1}{2} \sum_{x_1 \in S_1} \sum_{x_2 \in S_2} \left(\sqrt{p_{12}}(x_1, x_2) - \sqrt{p_1}(x_1) \sqrt{p_2}(x_2) \right)^2.$$

However, this measure is not ‘margin-free’. Recently, [Geenens \(2019\)](#) proposed a concept of copula suitable for discrete random vectors, in which the marginals are not made uniform by Probability Integral Transform but by the Iterated Proportional Fitting procedure (IPF). One can then apply (4.1) on \bar{p}_{12} , the ‘copula probability mass function’, that is, the joint distribution after the marginals have been made discrete uniform by IPF, making it now margin-free. That measure will be the topic of a follow-up paper, as well as the related axiomatic. Indeed, it is clear for instance that the concept of ‘pure dependence’ in (P6) only applies to continuous variables, and must be replaced by a different concept in the discrete case.

Supplementary Material

Reproducing scripts: Reproducing R-scripts for experiments contained in the article with descriptions included in files. (“HellingerReproduce.zip”)

Acknowledgement

This article includes results produced on the computational cluster Katana at UNSW Sydney.

Appendix

A Families of ϕ -divergences

Common choices for ϕ in (3.1) include:

(i) $\phi_{L,\alpha}(t) = |t - 1|^\alpha$, for $\alpha \geq 1$, yielding dependence measures

$D_{L,\alpha}(X_1, X_2) \propto \Delta_{\phi_{L,\alpha}}(F_{12} \| F_1 F_2)$ reminiscent of some sort of L_α -distance between F_{12} and $F_1 F_2$;

(ii) $\phi_{P,\alpha}(t) = (t^\alpha - \alpha t + \alpha - 1) / (\alpha(\alpha - 1))$, for $\alpha \in \mathbb{R} \setminus \{0, 1\}$, yielding dependence measures $D_{P,\alpha}(X_1, X_2) \propto \Delta_{\phi_{P,\alpha}}(F_{12} \| F_1 F_2)$ reminiscent of the ‘power divergence’ of [Cressie and Reid \(1984\)](#). As $\lim_{\alpha \rightarrow 1} \phi_{P,\alpha}(t) = t \log t - t + 1$, the Mutual Information $I(X_1, X_2) = \lim_{\alpha \rightarrow 1} D_{P,\alpha}(X_1, X_2)$ is a limiting case of this power-divergence family;

(iii) $\phi_{M,\alpha}(t) = |t^\alpha - 1|^{1/\alpha}$ for $\alpha \in (0,1]$, yielding dependence measures

$D_{M,\alpha}(X_1, X_2) \propto \Delta_{\phi_{M,\alpha}}(F_{12} \| F_1 F_2)$, reminiscent of Matusita (1967)'s divergence.

Now, in case (i), it is seen that for $\alpha = 1$, $\bar{\varphi} = 2 < \infty$, hence the rescaled measure $D_{L,1}^{\hat{\alpha}}(X_1, X_2) = (1/2)D_{L,1}(X_1, X_2)$ satisfies (P1)–(P7). By contrast, if $\alpha > 1$, then $\bar{\varphi} = \infty$, and the corresponding measure is unable to characterise pure dependence. That is the case of Pearson's Mean Square Contingency coefficient Φ^2 , which corresponds to $\alpha = 2$. In case (ii), one can check that, for $\alpha < 1$ (and $\alpha \neq 0$), $\bar{\varphi} = 1/(\alpha(1-\alpha)) < \infty$. Hence the associated dependence measure satisfies (P1)–(P7). For $\alpha > 1$, $\bar{\varphi} = \infty$ and one faces the same issue as above. That includes the limiting case $\alpha \rightarrow 1$, which yields the Mutual Information. Finally, in case (iii), $\bar{\varphi} = 2 < \infty$ for all $\alpha \in (0,1]$, making any 'Matusita' dependence measure $D_{M,\alpha}^{\hat{\alpha}}(X_1, X_2) = (1/2)D_{M,\alpha}(X_1, X_2)$ comply with (P1)–(P7).

All in all, among the ϕ -measures (3.1) of type (i)–(ii)–(iii), only $D_{L,1}^{\hat{\alpha}}(X_1, X_2)$, $D_{P,\alpha}^{\hat{\alpha}}$ for $\alpha < 1$ (and $\alpha \neq 0$) and $D_{M,\alpha}^{\hat{\alpha}}$ for $0 < \alpha \leq 1$ satisfy (P1)–(P7). Evidently, $D_{M,1}^{\hat{\alpha}}(X_1, X_2) = D_{L,1}^{\hat{\alpha}}(X_1, X_2)$. This particular measure, forming a kind of intersection between the L - and M -families, is Silvey (1964)'s Δ described in Section 3.2. At the intersection of the P - and M -families lies $D_{P,1/2}^{\hat{\alpha}}(X_1, X_2) = D_{M,1/2}^{\hat{\alpha}}(X_1, X_2)$, which is H^2 given in (4.1).

B Marginal transformations

Lemma 1 in Aya Moreno et al (2018) and Theorems 1 and 2 in Ebner et al (2018) establish the root- n consistency and the asymptotic normality of the oracle estimator B . However, those results hold only if the copula density c_{12} is bounded and bounded away from 0 on I^2 . This is a very restrictive assumption. In particular, many common parametric copula densities would grow unbounded in some of the corners of I^2 in the presence of dependence. Now, define a double marginal transformation

$\mathbf{T}_i = (T_{i1}, T_{i2}) = (\Xi_1^{-1}(U_{1i}), \Xi_2^{-1}(U_{2i}))$ where for $k = 1, 2$, Ξ_k^{-1} is the quantile function of a continuous distribution Ξ_k having a bounded density ξ_k on I . Standard

developments show that $\{\mathbf{T}_1, \dots, \mathbf{T}_n\}$ is a sample from a distribution having density $d_{12}(t_1, t_2) = c_{12}(\Xi_1(t_1), \Xi_2(t_2))\xi_1(t_1)\xi_2(t_2)$ on $(t_1, t_2) \in \text{Supp}(\Xi) \square \text{Supp}(\Xi_1) \times \text{Supp}(\Xi_2)$. Then we see that

$$B = \iint_{I^2} \sqrt{c_{12}}(u_1, u_2) du_1 du_2 = \iint_{\text{Supp}(\Xi)} \sqrt{d_{12}}(t_1, t_2) \sqrt{\xi_1}(t_1) \sqrt{\xi_2}(t_2) dt_1 dt_2,$$

by the obvious change-of-variable. Let $S_i = \min_{j \neq i} \|\mathbf{T}_j - \mathbf{T}_i\|_2$. Then, similarly to B , one can also estimate B by

$$B = \frac{2\sqrt{n-1}}{n} \sum_{i=1}^n S_i \sqrt{\xi_1}(T_{i1}) \sqrt{\xi_2}(T_{i2}).$$

The known weight function $\sqrt{\xi_1} \sqrt{\xi_2}$ is easily accounted for in the theoretical developments; see [Aya Moreno et al \(2018, Lemma 1\)](#) or [Ebner et al \(2018, Theorem 1\)](#).

The transformations Ξ_1 and Ξ_2 can be whatever is convenient. In [Geenens et al \(2017\)](#), their role was essentially to send the boundaries ‘far away’ from the observations so as to annihilate boundary effects. Here one could take, for instance, Ξ_1 and Ξ_2 to be Beta(6, 6)-distributions: the ensuing density d_{12} would remain supported on I^2 , but concentrated around the centre of it due to its Beta(6, 6)-marginals. That ‘double Beta transformation’ is also beneficial in terms of relaxing the above mentioned assumption on c_{12} . Indeed, let c_{12} satisfy Assumption 3.3 in [Geenens et al \(2017\)](#), which is rather mild and allows c_{12} to grow unboundedly in some of the corners of I^2 . With $\xi_k(t_k) \propto t_k^5(1-t_k)^5$ and $\Xi_k(t_k) = \int_0^{t_k} \xi_k(t^*) dt^*$, for $k = 1, 2$, that is the Beta(6, 6) density and cumulative distribution functions, it can be shown that d_{12} is Hölder continuous (with exponent $\alpha = 2$) on I^2 – this follows as in Lemma A.1 in [Geenens et al \(2017\)](#).

Then, Corollary 8 of [Singh and Póczos \(2016\)](#) applies, and B is root- n consistent for B . The results shown in Figures 5.1 and 5.2 were actually obtained making use of the Beta(6, 6) transformation.

C Cross-validation

For nonparametric functional estimation through orthogonal series approximation, it is well-known that the truncation cutoff plays the role of smoothing parameter (Efromovich, 1999, Chapter 2). Hence, if one wants to estimate \sqrt{c}_{12} by

$$\sqrt{c}_{12}(u_1, u_2) = \sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell} b_k(u_1) b_\ell(u_2),$$

as suggested in Section 5.3, one should select K and L appropriately. First, note that

$$\iint_{I^2} \left(\sqrt{c}_{12} \right)^2(u_1, u_2) du_1 du_2 = \sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^2,$$

which suggests to directly opt for the re-normalised estimator

$$\sqrt{c}_{12}^{\hat{a}}(u_1, u_2) = \frac{\sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell} b_k(u_1) b_\ell(u_2)}{\sqrt{\sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^2}}.$$

Obviously, $\iint_{I^2} \left(\sqrt{c}_{12}^{\hat{a}} \right)^2(u_1, u_2) du_1 du_2 \equiv 1$, for all $(K, L) \in \mathbb{N} \times \mathbb{N}$. The usual Mean Integrated Squared Error of this estimator is

$$\iint_{I^2} \left(\sqrt{c}_{12}^{\hat{a}}(u_1, u_2) - \sqrt{c}_{12}(u_1, u_2) \right)^2 du_1 du_2 = 2 - 2 \iint_{I^2} \sqrt{c}_{12}^{\hat{a}}(u_1, u_2) \sqrt{c}_{12}(u_1, u_2) du_1 du_2,$$

that we may seek to minimise with respect to K and L . It follows from Aya Moreno et al (2018, Lemma 1) that, for any bounded function

$f : I^2 \rightarrow \mathbb{R}$, $\iint_{I^2} f(u_1, u_2) \sqrt{c}_{12}(u_1, u_2) du_1 du_2$ can be estimated by:

$$\frac{2\sqrt{n-1}}{n} \sum_{i=1}^n \hat{R}_i f(\hat{U}_{i1}, \hat{U}_{i2});$$

compare (5.2). This justifies to estimate $\iint_{I^2} \sqrt{c}_{12}^{\hat{a}}(u_1, u_2) \sqrt{c}_{12}(u_1, u_2) du_1 du_2$ by

$$\frac{2\sqrt{n-1}}{n} \sum_{i=1}^n \hat{R}_i \sqrt{c_{12}}^{\hat{a}(-i)} (\hat{U}_{i1}, \hat{U}_{i2}) \quad (C.1)$$

where, as usual, the Leave-one-Out version of the estimator $\sqrt{c_{12}}^{\hat{a}(-i)}$ is used for avoiding overfitting. Explicitly, this is

$$\sqrt{c_{12}}^{\hat{a}(-i)}(u_1, u_2) = \frac{\sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^{(-i)} b_k(u_1) b_\ell(u_2)}{\sqrt{\sum_{k=0}^K \sum_{\ell=0}^L \{\hat{\beta}_{k\ell}^{(-i)}\}^2}}$$

where

$$\hat{\beta}_{k\ell}^{(-i)} = \frac{2\sqrt{n-2}}{n-1} \sum_{i' \neq i} \hat{R}_{i'}^{(-i)} b_k(\hat{U}_{i'1}) b_\ell(\hat{U}_{i'2})$$

and $\hat{R}_{i'}^{(-i)} = \min_{j \notin \{i, i'\}} \|\mathbf{U}_j - \mathbf{U}_{i'}\|_2$. Note that, for any i , $\hat{R}_{i'}^{(-i)} \neq \hat{R}_{i'}$ only if \mathbf{U}_i is the closest neighbour of $\mathbf{U}_{i'}$, which concerns only a small number of $i' \in \{1, \dots, n\}$; in such a case $\hat{R}_{i'}^{(-i)}$ is the distance from $\mathbf{U}_{i'}$ to its second closest neighbour. The values $\hat{R}_{i'}^{(-i)}$ are thus very easily obtained, and so are the $\hat{\beta}_{k\ell}^{(-i)}$'s and eventually $\sqrt{c_{12}}^{\hat{a}(-i)}(u_1, u_2)$. Finally, K and L may be chosen as the values which maximise (C.1), the (estimated) Battacharyya affinity coefficient between c_{12} and its estimator, that is

$$\frac{2\sqrt{n-1}}{n} \sum_{i=1}^n \hat{R}_i \frac{\sum_{k=0}^K \sum_{\ell=0}^L \hat{\beta}_{k\ell}^{(-i)} b_k(\hat{U}_{i1}) b_\ell(\hat{U}_{i2})}{\sqrt{\sum_{k=0}^K \sum_{\ell=0}^L \{\hat{\beta}_{k\ell}^{(-i)}\}^2}}.$$

These values are easy to identify given that K and L are integers.

D Empirical estimation and power study

The aim of this section is to investigate the ability of the empirical Hellinger correlation $\hat{\eta}$ to consistently estimates its population counterpart η and to detect (i.e., to test) the dependence between two univariate random variables X_1 and X_2 . All our computing codes are available in Supplementary Material.

They have been written using the version 3.6.1 of the \mathbb{R} programming language (R Core Team, 2019), along with parts written in C++. We used version 1.3 of our \mathbb{R} package `HellCor`. Simulations have been performed on a computational cluster using a 64-bit CPU and 8GB of RAM. Long computation times are indicated at appropriate places in the next sections.

D.1 A unifying Data Generating Process

We consider a unifying data generating process (DGP) for the vector (X_1, X_2) , which can be written under the following form:

$$\begin{pmatrix} X_1 \\ X_2 \end{pmatrix} = \begin{pmatrix} \Psi_1(U) \\ \alpha\Psi_2(U) + (1-\alpha)\Psi_2(V) \end{pmatrix} + \beta\delta, \quad (\text{D.1})$$

where U and V are independent $U_{[0,1]}$ -random variables, δ is an independent bivariate standard normal noise (with noise level set via the parameter $\beta \geq 0$) and Ψ_1 and Ψ_2 are ‘shape’ functions from $[0,1]$ to \mathbb{R} . The parameter $\alpha \in [0,1]$ determines the strength of the dependence between X_1 and X_2 . This is somehow related to the degree of freedom discussed in Section 2.3. When $\alpha = 0$, X_1 and X_2 are independent. When $\alpha = 1$ and $\beta = 0$, they are purely dependent. The functions Ψ_1 and Ψ_2 characterise the joint distribution of (X_1, X_2) . The ones we consider in this empirical study are explicitly given in Table D.1, in a unifying form that makes it easy to simulate from (see the undocumented functions `datagen15()` and `gendep()` in our \mathbb{R} package `HellCor`).

The marginal distribution of X_1 is unrelated to the choice of α in $[0,1]$, whereas for X_2 it remains unchanged only when $\alpha \in \{0,1\}$ (i.e., for the two extreme cases). The scenarios of Table D.1 are directly inspired from those considered in Heller et al (2012, Figure 4) – also, our Figure 5.2. However, playing with the parameters α and β , we can create from each of them a whole family of distributions showing any level of dependence and/or any level of noise. Figures D.1, D.2 and D.3 show typical scatter-plots of samples of size $n = 500$ drawn from those 15 scenarios, for pure dependence ($\alpha = 1$ and $\beta = 0$), mild dependence ($\alpha = 0.7$ and $\beta = 0.01$) and independence ($\alpha = 0$ and $\beta = 0.1$),

respectively. The values of $\hat{\eta}$ are added on top of each one of these plots. One can definitely relate the value of $\hat{\eta}$ to the strength of the dependence apparent on the plots.

D.2 Accuracy of estimation

Here we explore the ability of our measure to consistently estimate the true population value of the Hellinger correlation for the 15 situations shown in Figure D.2, that is, with $\alpha = 0.7$ and $\beta = 0.01$ in (D.1). Our graphical findings are displayed in Figures D.4–D.18, for samples of size $n = 100$ to $n = 2,000$. The left panel (I) shows the boxplots of the estimated values of η , while the right panels (II) and (III) show the values of K and L returned by the cross-validation criterion for enforcing the normalisation (Section 5.3). (Processing time for each figure is around 2,000s.) As noticed in Section 5.4, the real value of η is not available in these situations, so we have approximated η by $\hat{\eta}$ computed on a sample of size $n = 20,000$. The Mean Squared Error (shown on top of each violin plot, panel (I)) quickly decreases with the sample size for all scenarios. Also, the selected values of K and L tend to increase with the sample size, as we would expect for smoothing parameters in a nonparametric density estimation context.

Now, two cases require larger samples for providing accurate estimation of η , though: ‘Spiral’ and ‘4-Circles’ (Figures D.14–D.15). This can easily be understood from Figure D.2. If dependence is obvious visually on the other scatter-plots (and quantified by a high value of $\hat{\eta}$), it is not at all the case for those two scenarios. Indeed it is actually very difficult to discern whether the data show just noise or not. The empirical Hellinger correlation is just significant (at level $\alpha = 0.05$) for ‘Spiral’, and not significant for ‘4-Circles’ – remind from Section 5.5 that the critical value at that level for $n = 500$ is $\bar{\eta}_c = 0.146$. Given the nature of the relationship, and the level of noise involved, it would take a larger sample than $n = 500$ to be able to grab the dependence with confidence. The effect of the normalisation is also interesting to observe. For small samples, when the data look really like noise, the cross-validation criterion consistently picks $K = L = 1$ – see Figures D.14–

D.15, panels (II) and (III). The procedure attempts to fit a copula density as flat as possible, and the normalisation tends to hold $\hat{\eta}$ down to 0. For larger sample sizes, when the dependence becomes perceivable, the CV-criterion picks higher values for K and L , allowing $\hat{\eta}$ to grow. The switch between these two ‘regimes’ happens at around $n = 500$ for ‘Spiral’, and around $n = 1,000$ for ‘4-Circles’. Around those sample sizes, the distribution of $\hat{\eta}$ is essentially bi-modal, depending on the values of K and L selected by cross-validation, that is, depending on whether the dependence can be perceived or not. This is actually in agreement with intuition.

D.3 Power study

Finally, we complement here the simulation study shown in Section 5.5 to compare the empirical power of independence tests based on $\hat{\eta}_0$ and $\hat{\eta}$ (non-normalised and normalised version of the empirical Hellinger correlation) to some well known competitors, namely dCov: distance correlation ([Székely et al, 2007](#)); HHG: [Heller et al \(2012\)](#); \mathcal{G}^2 : [Wang et al \(2017\)](#); RCD: robust copula dependence ([Ding et al, 2017](#)); GMR: [Granger et al \(2004\)](#). For the GMR test, we used the implementation of ([Maasoumi and Racine, 2002](#)) provided in the `np R` package ([Hayfield and Racine, 2008](#)) available on the CRAN. For the \mathcal{G}^2 test of [Wang et al \(2017\)](#), we used the implementation provided in the `Gs R` package.² We reprogrammed the other tests in C++. Tables D.2 and D.3 summarise the empirical power results, for sample sizes $n = 20$ and $n = 50$ respectively, based on $M = 10,000$ Monte Carlo replications of each scenario.

It appears that the test based on $\hat{\eta}$ exhibits a good power overall when the sample size reaches $n = 50$, and is ranked among the three best tests of independence (with HHG and RCD). Compared to the test based on $\hat{\eta}_0$, it is seen that the normalisation is not always detrimental to the power of detecting dependence.

References

Ali, M.S. and Silvey, S.D. (1965), Association between random variables and the dispersion of a Radon-Nikodym derivative, J. R. Stat. Soc. Ser. B. Stat. Methodol., 27, 100-107.

Ali, M.S. and Silvey, S.D. (1966), A general class of coefficients of divergence of one distribution from another, J. R. Stat. Soc. Ser. B. Stat. Methodol., 28, 131-140.

Arel-Bundock, V. and Enevoldsen, N. and Yetman, C.-J. (2018), countrycode: An R package to convert country names and country codes, J. Open Source Softw., 3, 848.

Aya Moreno, C., Geenens, G. and Penev, S. (2018), Shape-preserving wavelet-based multivariate density estimation, J. Multivariate Anal., 168, 30-47.

Balakrishnan, N. and Lai, C.D., Continuous Bivariate Distributions, Springer, 2009.

Beare, B. (2010), Copulas and temporal dependence, Econometrica 78, 395-410.

Bell, C.B. (1962), Mutual information and maximal correlation as measures of dependence, Ann. Math. Statist., 33, 587-595.

Bergsma, W. and Dassios, A. (2014), A consistent test of independence based on a sign covariance related to Kendall's tau, Bernoulli, 20, 1006-1028.

Berrett, T.B., Samworth, R.J. and Yuan, M. (2018), Efficient multivariate entropy estimation via k -nearest neighbour distances, Ann. Statist., 47, 1, 288-318.

Berrett, T.B. and Samworth, R.J. (2019), Nonparametric independence testing via mutual information, Biometrika, 106, 547-566.

Bhattacharyya, A. (1943), On a measure of divergence between two statistical populations defined by their probability distributions, Bull. Calcutta Math. Soc., 35, 99-109.

Boglioni-Beaulieu, G. (2016), A consistent test of independence between random vectors, MSc thesis, Université de Montréal.

Cover, T.M. and Thomas, J.A., Elements of Information Theory, 2nd Edition, Wiley, New York, 2006.

Cressie, N. and Reid, T.R.C. (1984), Multinomial goodness-of-fit tests, J. R. Stat. Soc. Ser. B. Stat. Methodol., 46, 440-464.

Deheuvels, P. (2009), A multivariate Bahadur-Kiefer representation for the empirical copula process, J. Math. Sci. (N.Y.), 163, 382-398.

Dette, H., Siburg, K.F. and Stoimenov, P.A. (2013), A Copula-Based Non-parametric Measure of Regression Dependence, Scand. J. Statist., 40, 21-41.

Ding, A.A., Dy, J.G., Li, Y. and Chang, Y. (2017), A robust-equitable measure for feature ranking and selection, J. Mach. Learn. Res., 18, 1-46.

Durante, F., Fernández-Sánchez, J. and Sempì, C. (2013), A note of notion of singular copula, Fuzzy Sets and Systems, 211, 120-122.

Ebner, B., Henze, N. and Yukich, J.E. (2018), Multivariate goodness-of-fit on flat and curved spaces via nearest neighbor distances, J. Multivariate Anal., 165, 231-242.

Efromovich, S., Nonparametric curve estimation: methods, theory and applications, Springer-Verlag, New-York, 1999.

Geenens, G., Charpentier, A. and Paindaveine, D. (2017), Probit transformation for nonparametric kernel estimation of the copula density, Bernoulli, 23, 1848-1873.

Geenens, G. (2019), An essay on copula modelling for discrete random vectors; or how to pour new wine into old bottles, *Manuscript*, arXiv:1901.08741.

Genest, C. and Neslehová, J. (2007), A primer on copulas for count data, *ASTIN Bull.*, 37, 475-515.

Graham, N.A.J., Wilson, S.K., Carr, P., Hoey, A.S., Jennings, S. and MacNeil, M.A. (2018), Seabirds enhance coral reef productivity and functioning in the absence of invasive rats, *Nature*, 559, 250-253.

Granger, C.W., Maasoumi, E. and Racine, J. (2004), A dependence metric for possibly nonlinear processes, *J. Time Series Anal.*, 25, 649-669.

Gretton, A., Bosquet, O., Smola, A. and Schölkopf, B. (2012), Measuring statistical dependence with Hilbert-Schmidt norms, *In: Algorithmic Learning Theory*, Jain, S., Simon, H.U. and Tomita, E. Eds, Springer, pp. 63-77.

Hall, P. (1992), The bootstrap and Edgeworth expansions, *Springer Series in Statistics*, Springer-Verlag, New-York.

Hayfield, T. and Racine, J. S. (2008), Nonparametric econometrics: The `np` package, *J. Stat. Softw.*, 27.

Heller, R., Heller, Y. and Gorfine, M. (2012), A consistent multivariate test of association based on ranks of distance, *Biometrika*, 100(2), 503-510.

Heller, R., Heller, Y., Kaufman, S., Brill, B. and Gorfine, M. (2016), Consistent distribution-free K sample and independence tests for univariate random variables, *J. Mach. Learn. Res.*, 17, 1-54.

Hoeffding, W. (1942), Stochastic dependence and functional relationships, *Skandinavisk Aktuarietidskrift*, 25, 200-227; Recovered from: *The Collected Works of Wassily Hoeffding*, Eds: Fisher, N.I. and Sen, P.K., *Springer Series in Statistics*, Springer Verlag, 1994, pp.135-155.

Hoeffding, W. (1948), A nonparametric test of independence, *Ann. Math. Statist.*, 19, 546-557.

Hong, Y. and White, H. (2005), Asymptotic distribution theory for nonparametric entropy measures of serial dependence, *Econometrica*, 73, 837-901.

Janzing, D., Balduzzi, D., Grosse-Wentrup, M. and Schölkopf, B. (2013), Quantifying causal influences, *Ann. Statist.*, 41, 2324-2358.

Karlsson, S. and Löthgren, M. (2000), Computationally efficient double bootstrap variance estimation, *Comput. Statist. Data Anal.*, 33, 237-247.

Kimeldorf, G. and Sampson, A.R. (1978), Monotone dependence, *Ann. Statist.*, 6, 895-903.

Kinney, J. and Atwal, G. (2014), Equitability, mutual information, and maximal information coefficient, *Proc. Natl. Acad. Sci. USA*, 111, 3354-3359.

Kiriliouk, A., Segers, J. and Tsukahara, H. (2019), On some resampling procedures with the empirical beta copula, *Manuscript*, arXiv:1905.12466.

Kozachenko, L. and Leonenko, N. (1987), On statistical estimation of entropy of a random vector, *Problems Inform. Transmission*, 23, 95-101.

Lancaster, H.O. (1963), Correlation and complete dependence of random variables, *Ann. Math. Statist.*, 34, 1315-1321.

Lancaster, H.O. (1982), Measures and indices of dependence, *In: Kotz, S. and Johnson, N.L. (eds), Encyclopedia of Statistical Sciences*, Vol. 2, Wiley, New York, pp. 334-339.

Leonenko, N., Pronzato, L. and Savani, V. (2008), A class of Rényi information estimators for multidimensional densities, *Ann. Statist.*, 36, 2153-2182.

Liese, F. and Vajda, I. (2006), On divergences and informations in statistics and information theory, *IEEE Trans. Inform. Theory*, 52, 4394-4412.

Liitiäinen, E., Lendasse, A. and Corona, F. (2010), A boundary corrected expansion of the moments of nearest neighbor distributions, *Random Struct. Algorithms*, 37, 223-247.

Linfoot, E.H. (1967), An informational measure of correlation, *Information and Control*, 1, 85-99.

Maasoumi, E. and Racine, J. (2002), Entropy and predictability of stock market returns, *J. Econometrics*, 107, 291-312.

Matusita, K. (1967), On the notion of affinity of several distributions and some of its applications, *Ann. Inst. Statist. Math.*, 19, 181-192.

Micheas, A.C. and Zografos, K. (2006), Measuring stochastic dependence using φ -divergence, *J. Multivariate Anal.*, 97, 765-784.

Mikusinski, P., Sherwood, H. and Taylor, M.D. (1992), Shuffles of Min, *Stochastica*, 13, 61-74.

Pfister, N., Bühlmann, P., Schölkopf, B. and Peters, J. (2018), Kernel-based tests for joint independence, *J. R. Stat. Soc. B*, 80, 5-31.

Póczos, B., Ghahramani, Z. and Schneider, J. (2012), Copula-based kernel dependency measures, *In: Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 775-782.

Pyke, R. (1965), Spacings, *J. R. Stat. Soc. B*, 27, 395-449.

R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

Reimherr, M. and Nicolae, D.L. (2013), On quantifying dependence: a framework for developing interpretable measures, *Statist. Sci.*, 28, 116-130.

Rényi, A. (1959), On measures of dependence, Acta. Math. Acad. Sci. Hungar., 10, 441-451.

Reshef, D., Reshef, Y., Finucane, H., Grossman, S., McVean, G., Turnbaugh, P., Lander, E., Mitzenmacher, M. and Sabeti, P. (2011), Detecting novel associations in large data sets, Science, 334, 1518-1524.

Reshef, Y., Reshef, D., Finucane, H., Sabeti, P. and Mitzenmacher, M. (2016), Measuring dependence powerfully and equitably, J. Mach. Learn. Res., 17, 1-63.

Rosenblatt, M. (1975), A quadratic measure of deviation of two-dimensional density estimates and a test of independence, Ann. Statist., 3, 1-14.

Schweizer, B. and Wolff, E. (1981), On nonparametric measures of dependence for random variables, Ann. Statist., 9, 879-885.

Segers, J., Sibuya, M. and Tsukahara, H. (2017), The empirical beta copula, J. Multivariate Anal., 155, 35-51.

Sejdinovic, D., Sriperumbudur, B., Gretton, A. and Fukumizu, K. (2013), Equivalence of distance-based and RKHS-based statistics in hypothesis testing, Ann. Statist., 41, 2263-2291.

Silvey, S.D. (1964), On a measure of association, Ann. Math. Statist., 35, 1157-1166.

Singh, S. and Póczos, B. (2016), Finite-sample analysis of fixed- k nearest neighbor density functional estimators, *In: Advances in Neural Information Processing Systems*, pp. 1217-1225.

Skaug, H. and Tjøstheim, D. (1996), Testing for serial independence using measures of distances between densities, *In: Athens Conference on Applied Probability and Time Series*, P. Robinson and M. Rosenblatt Eds., Springer Lecture Notes in Statistics.

Steele, J.M., Probability theory and combinatorial optimization, Vol. 69, Siam, 1997.

Su, L. and White, H. (2008), A nonparametric Hellinger metric test for conditional independence, *Econometric Theory*, 24, 829-864.

Sun, N. and Zhao, H. (2014), Putting things in order, *Proc. Natl. Acad. Sci. USA*, 111, 16236-16237.

Székely, G., Rizzo, M. and Bakirov, N. (2007), Measuring and testing dependence by correlation of distances, *Ann. Statist.*, 35, 2769-2794.

Székely, G. and Rizzo, M. (2009), Brownian distance covariance, *Ann. Appl. Stat.*, 3, 1236-1265.

Wang, X., Jiang, B. and Liu, J.S. (2017), Generalized R-squared for detecting dependence, *Biometrika*, 104, 129-139.

Weihs, L., Drton, M. and Meinshausen, N. (2018), Symmetric rank covariances: a generalized framework for nonparametric measures of dependence, *Biometrika*, 105, 547-562.

Zeng, X., Xia, Y. and Tong, H. (2018), Jackknife approach to the estimation of mutual information, *Proc. Natl. Acad. Sci. USA*, 115, 9956-9961.

Zhang, K. (2019), BET on independence, *J. Amer. Statist. Assoc.*, 114(528), 1620-1637.




Fig. 5.1 $M = 1,000$ replications of 6 estimators of η on bivariate Gaussian samples of size $n = 500$ with $\rho = 0.4$ (left) and $\rho = 0.8$ (right).

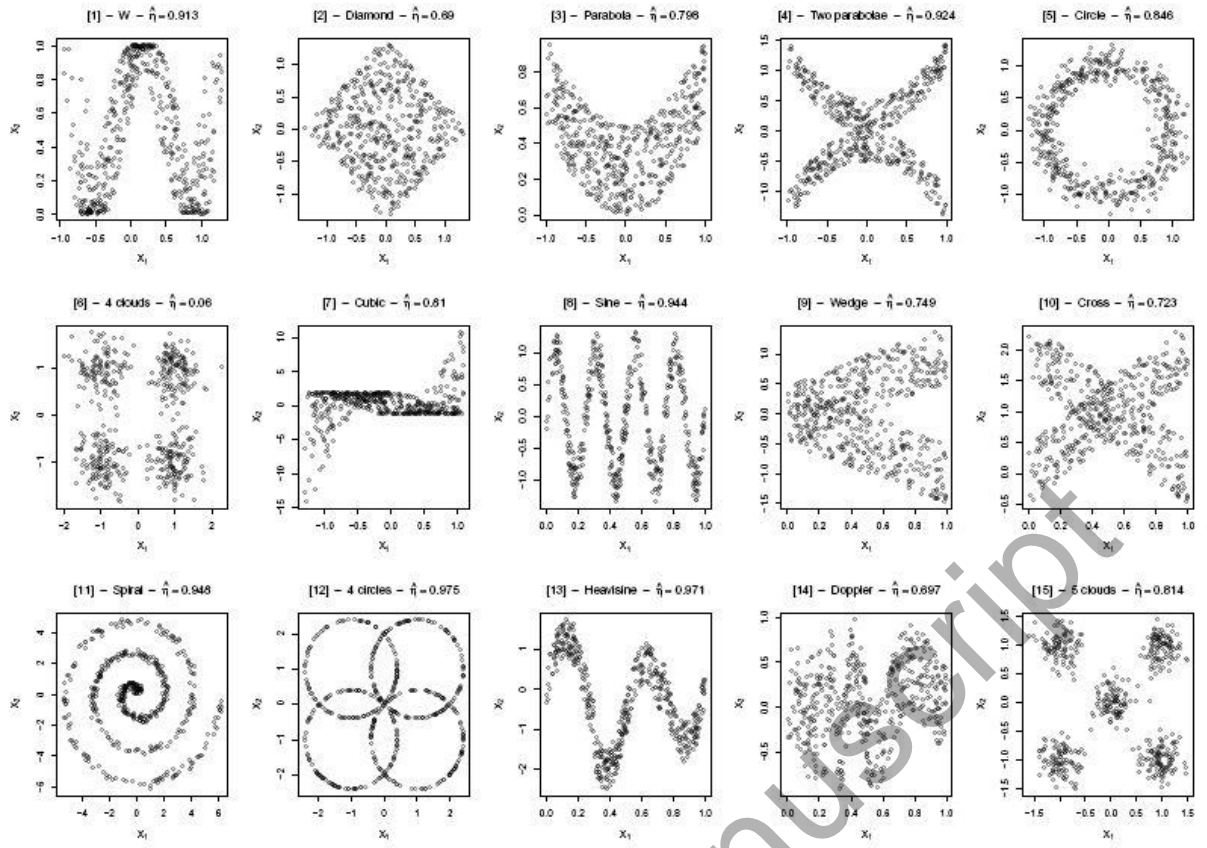


Fig. 5.2 Fifteen random samples of size $n = 500$ generated from the 15 scenarii of [Heller et al \(2016, Figure 4\)](#). The empirical Hellinger correlations are shown on top.

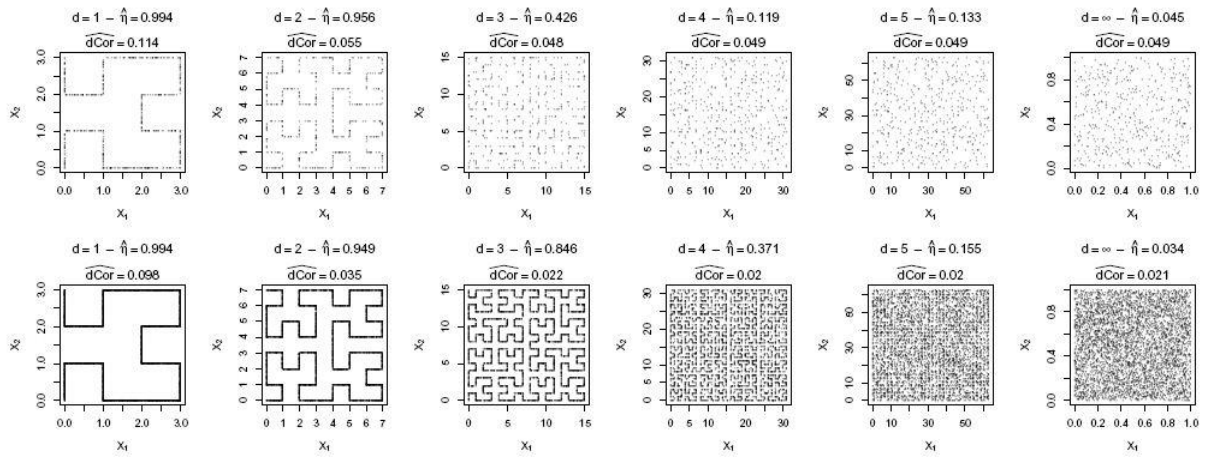


Fig. 5.3 Typical random samples of $n = 500$ (top row) and $n = 5,000$ (bottom row) points uniformly distributed on the Peano curve at resolution $d = 1, 2, 3, 4, 5$ and $d = \infty$ (independence). The empirical Hellinger correlations are shown on top, as well as the empirical distance correlations (\widehat{dCor}).

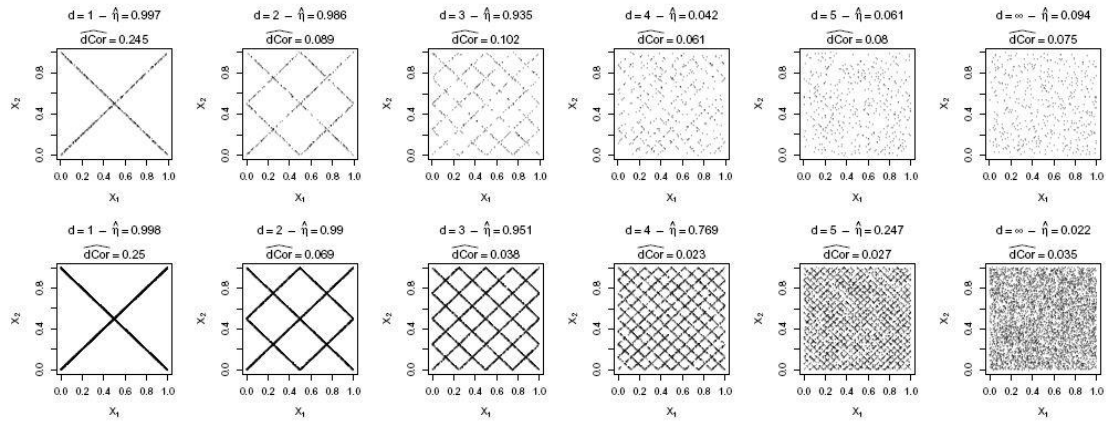


Fig. 5.4 Typical random samples of $n = 500$ (top row) and $n = 5,000$ (bottom row) points uniformly distributed on the 'bisection expanding cross' (Zhang, 2019) at resolution $d = 1, 2, 3, 4, 5$ and $d = \infty$ (independence). The empirical Hellinger correlations are shown on top, as well as the empirical distance correlations (\widehat{dCor}).

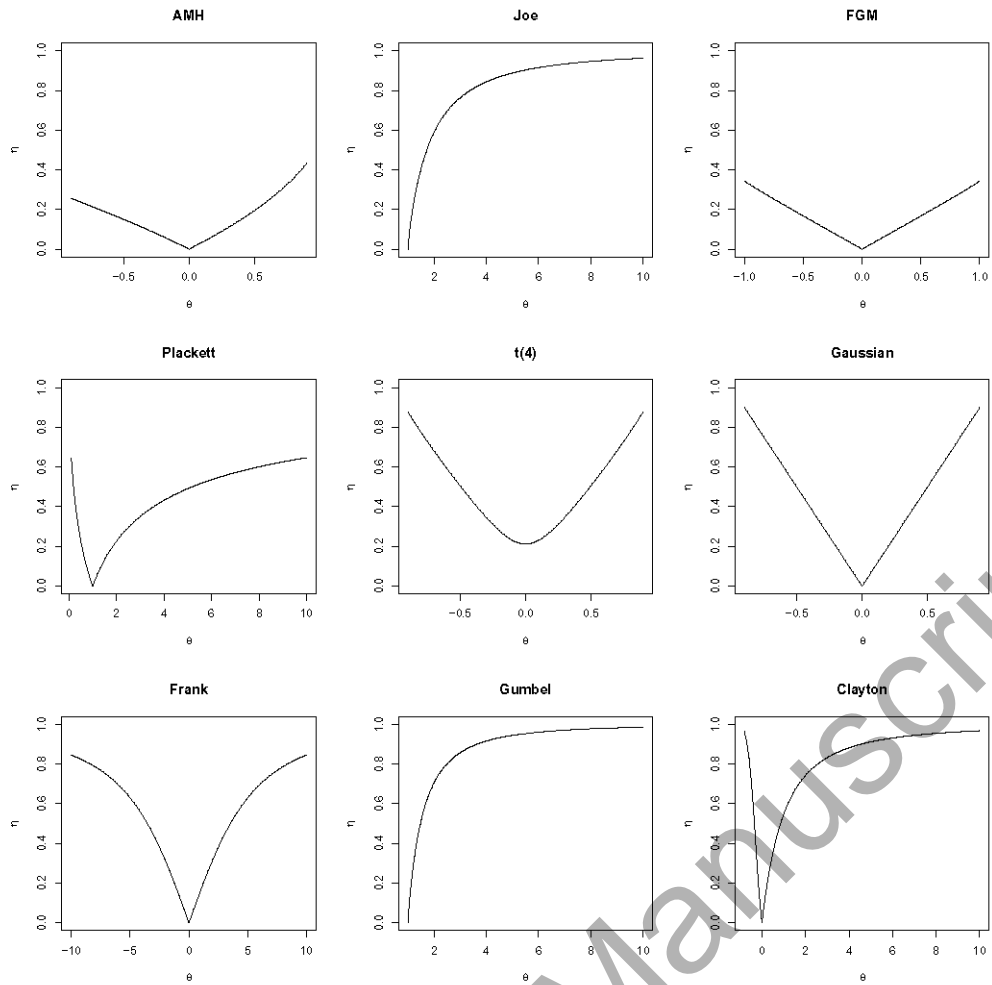


Fig. 5.5 The Hellinger correlation η as a function of the classical parameter θ for nine common copulas.

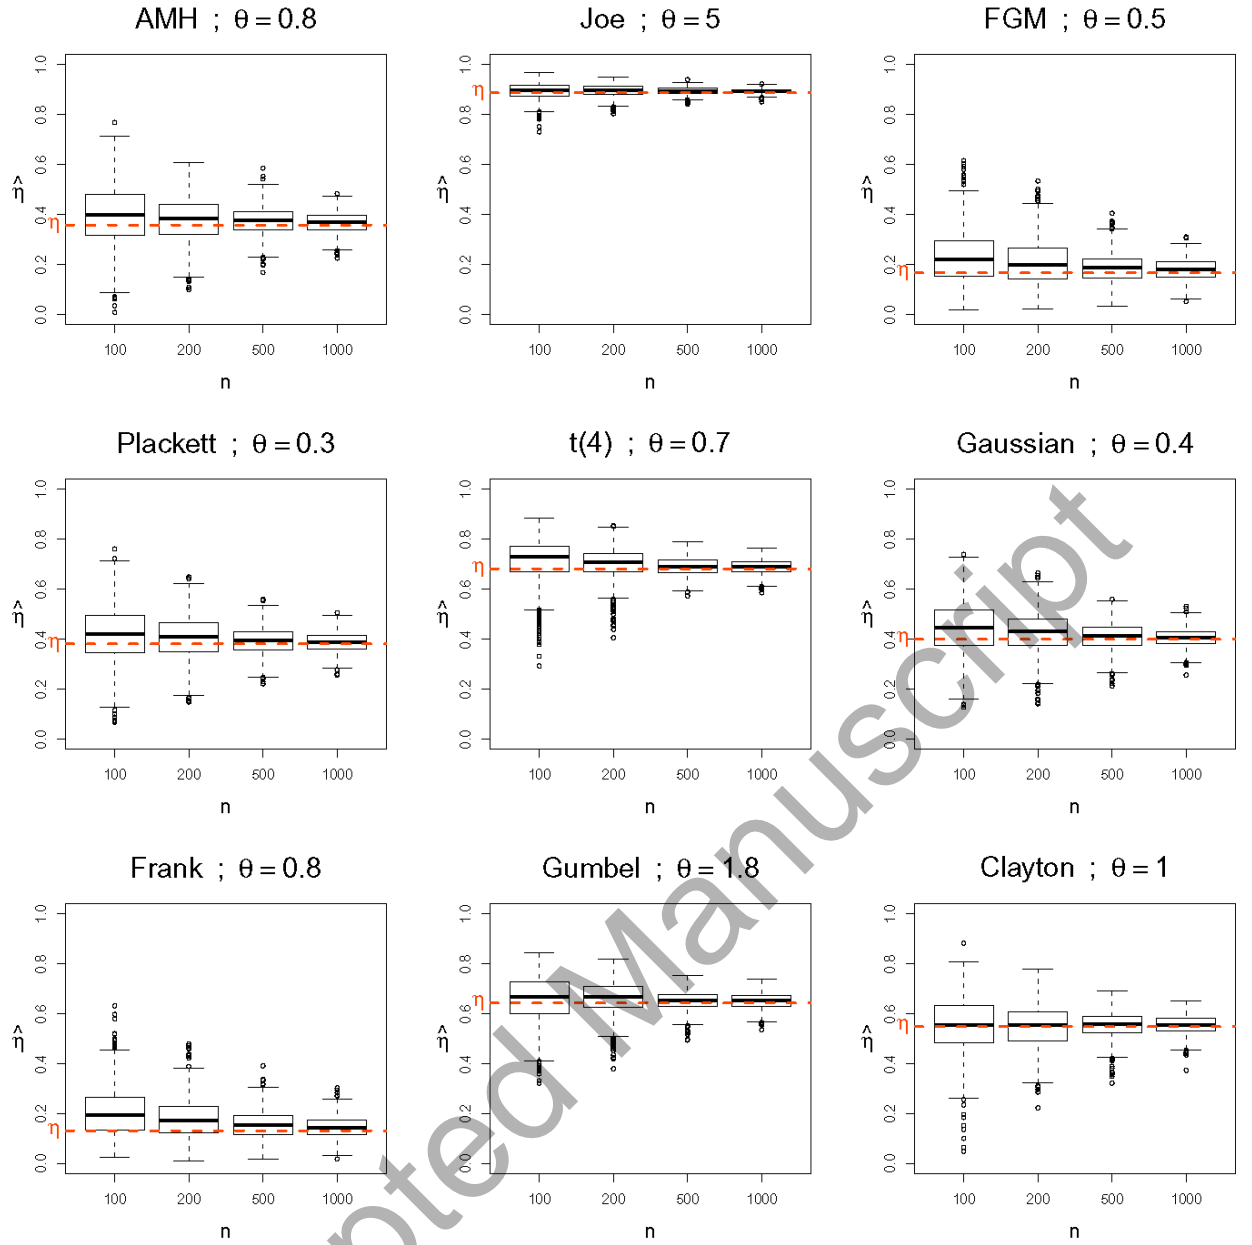


Fig. 5.6 $M = 1,000$ replications of $\hat{\eta}$ for each copula of Table 5.1 and each sample size $n = 100, 250, 500$ and $1,000$. The corresponding true values of η are indicated by the horizontal dashed line.

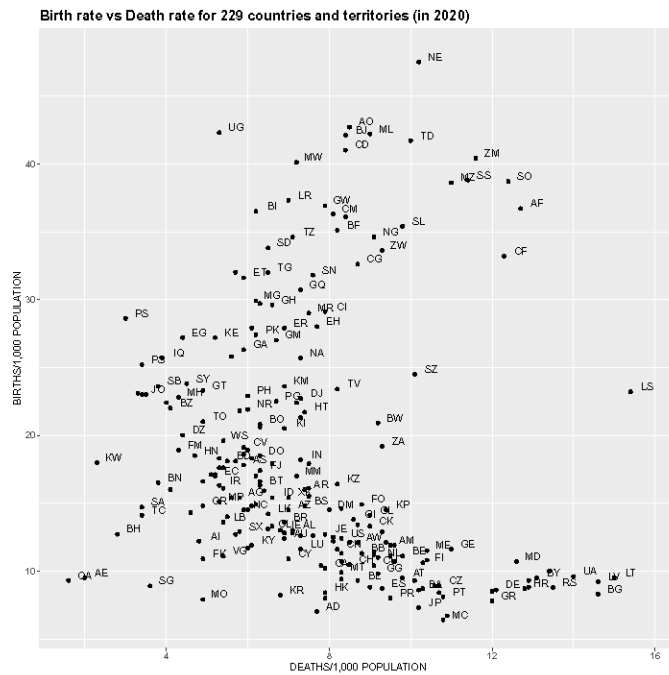


Fig. 6.1 Birth rate versus death rate in 229 countries and island territories around the world in 2020. (Source: The CIA World Fact Book.) These data are made available in our `HellCor` package, via the command `data(worlddemographics)`. Countries are represented by their 2-digit ISO code (obtained from the R package `countrycode` (Arel-Bundock et al, 2018) or extracted from the website countrycode.org).

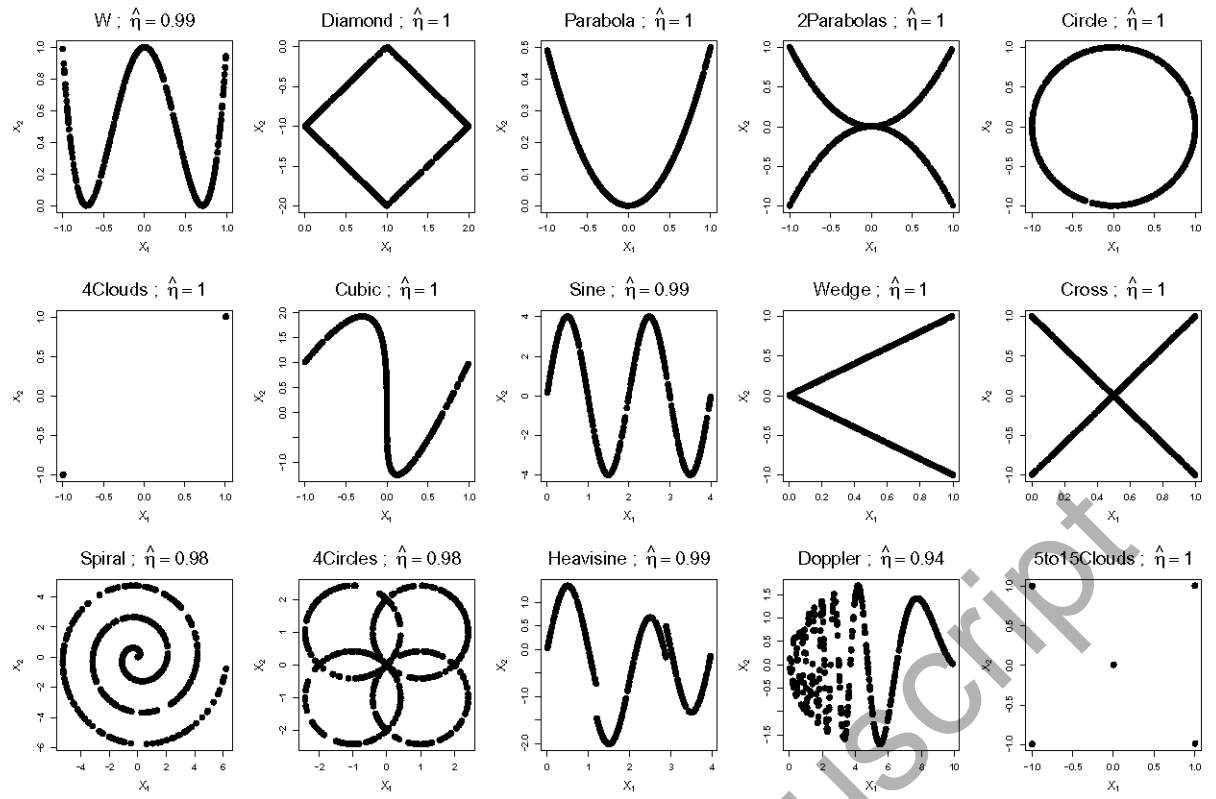


Fig. D.1 Pure dependence: $\alpha = 1$, $\beta = 0$; $n = 500$.

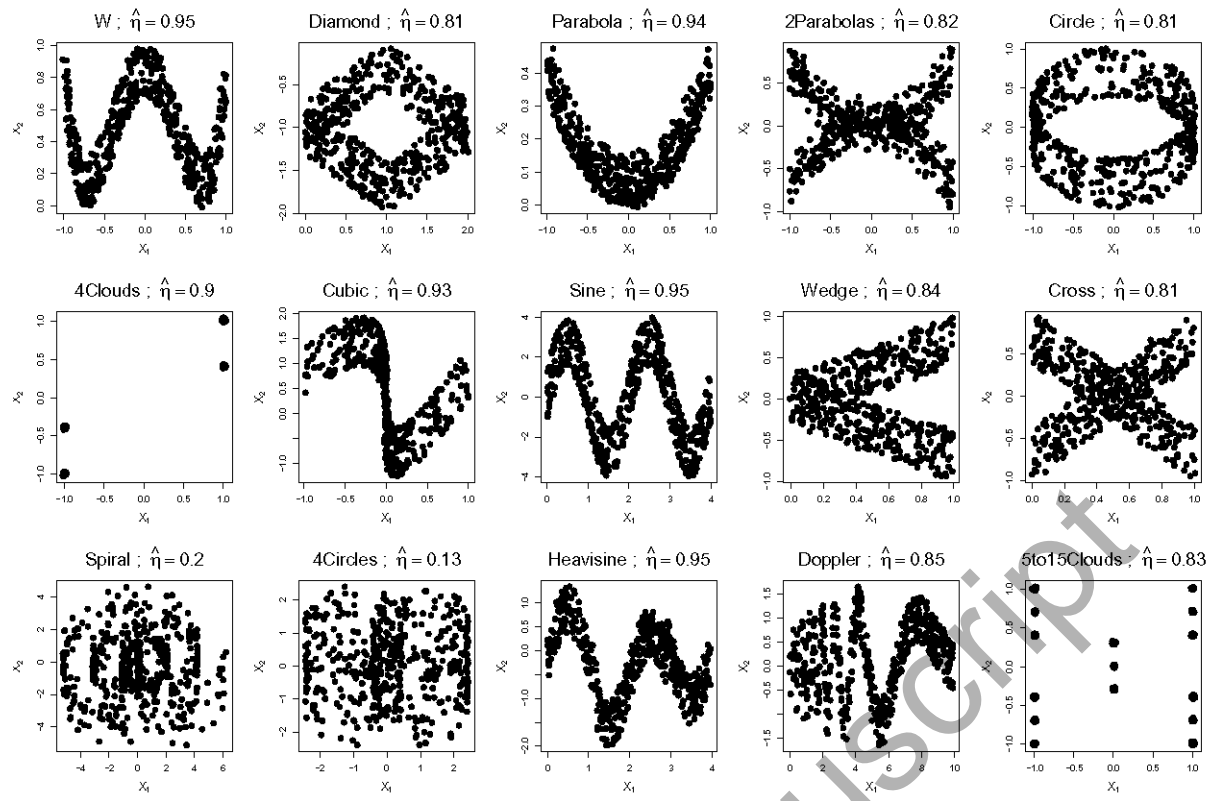


Fig. D.2 Dependence: $\alpha = 0.7, \beta = 0.01; n = 500$

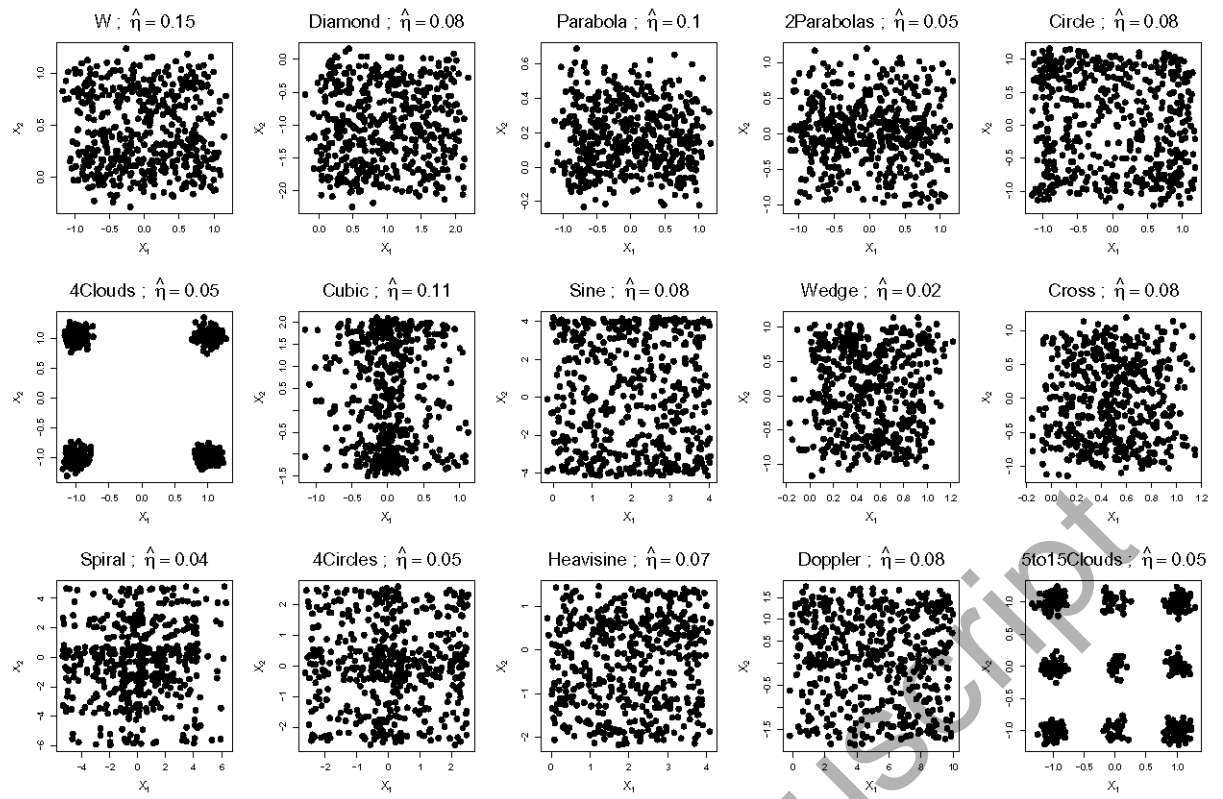




Fig. D.3 Independence: $\alpha = 0, \beta = 0.1$; $n = 500$.

I



II



III




Fig. D.4 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$.

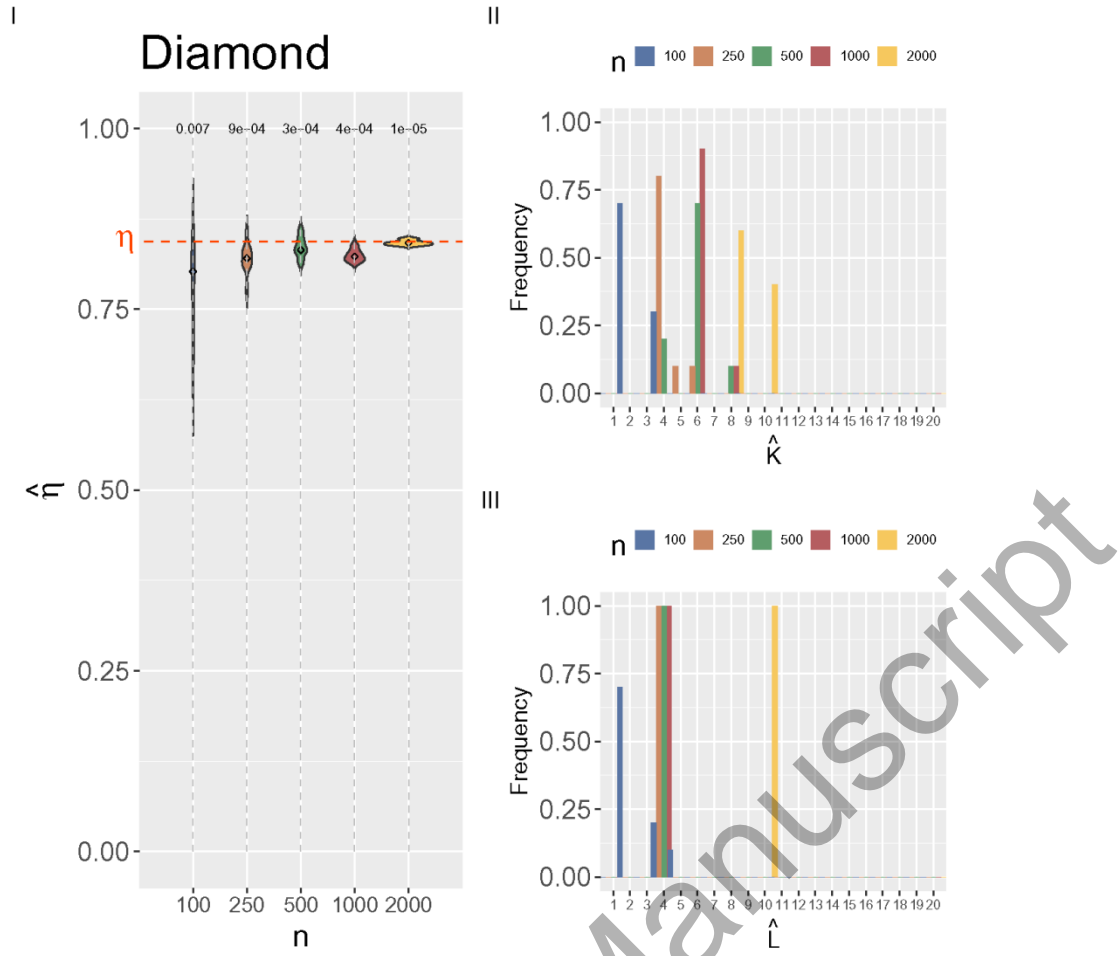




Fig. D.5 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III

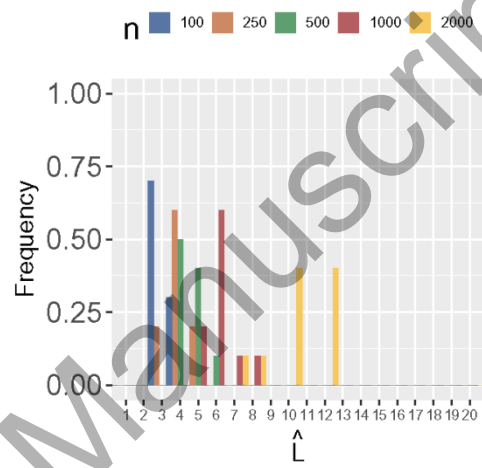


Fig. D.6 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

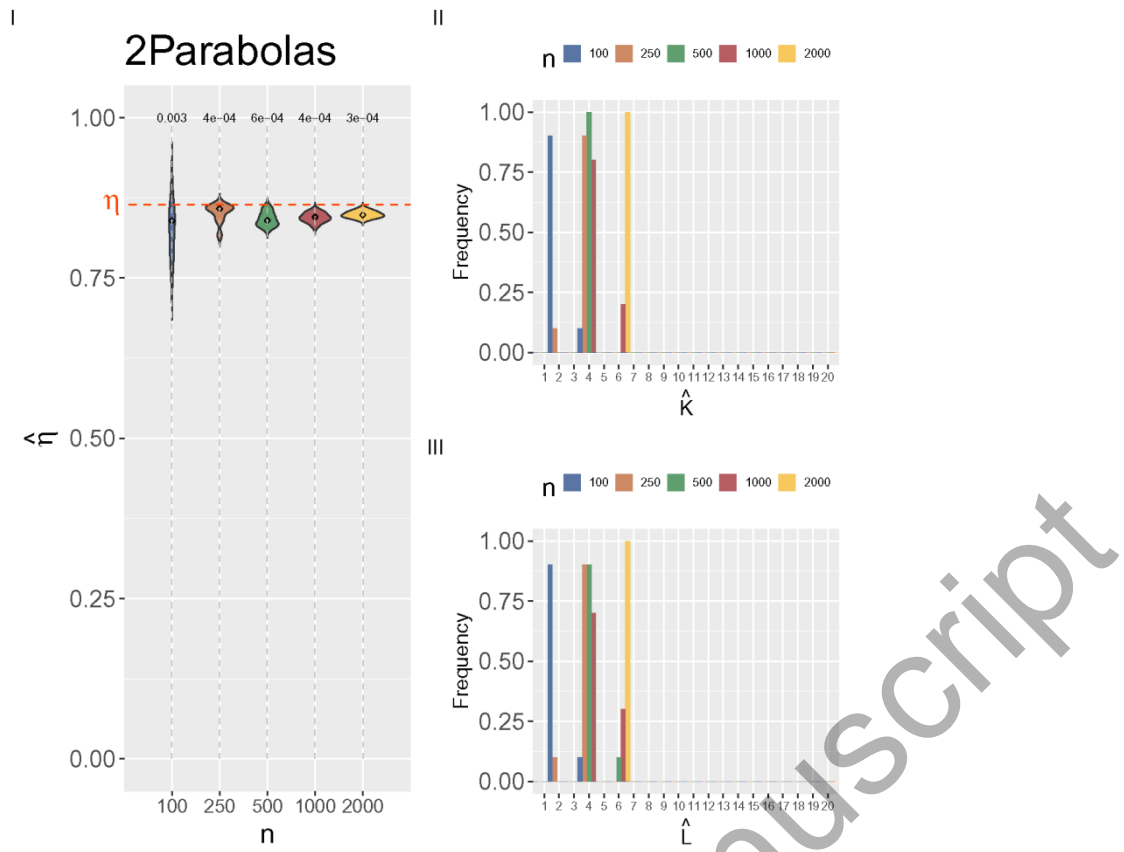
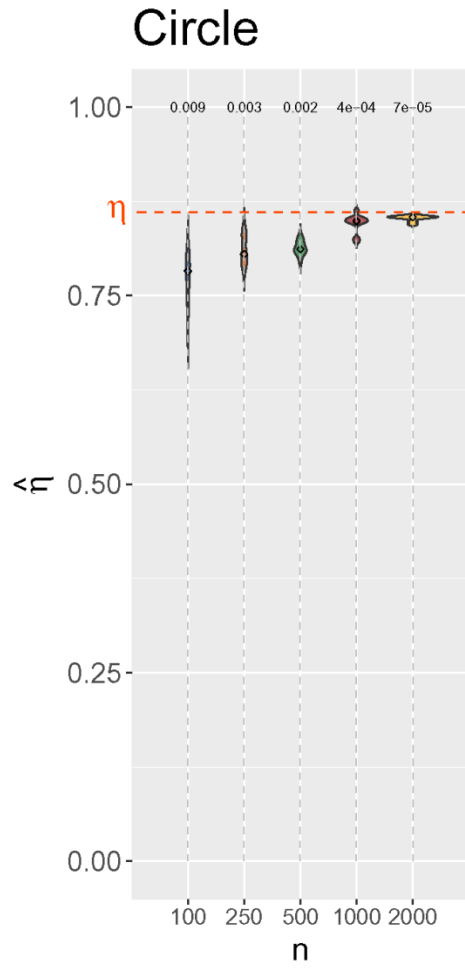



Fig. D.7 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III



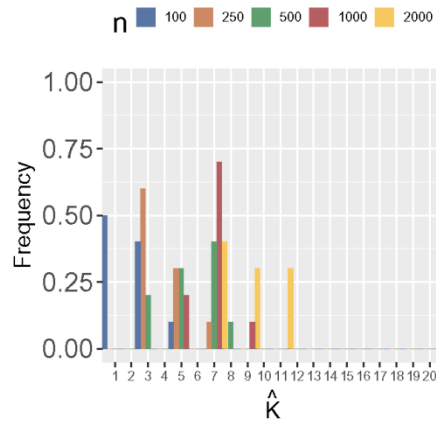


Fig. D.8 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III

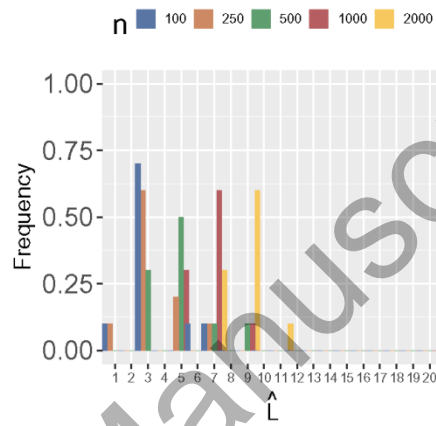




Fig. D.9 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III




Fig. D.10 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

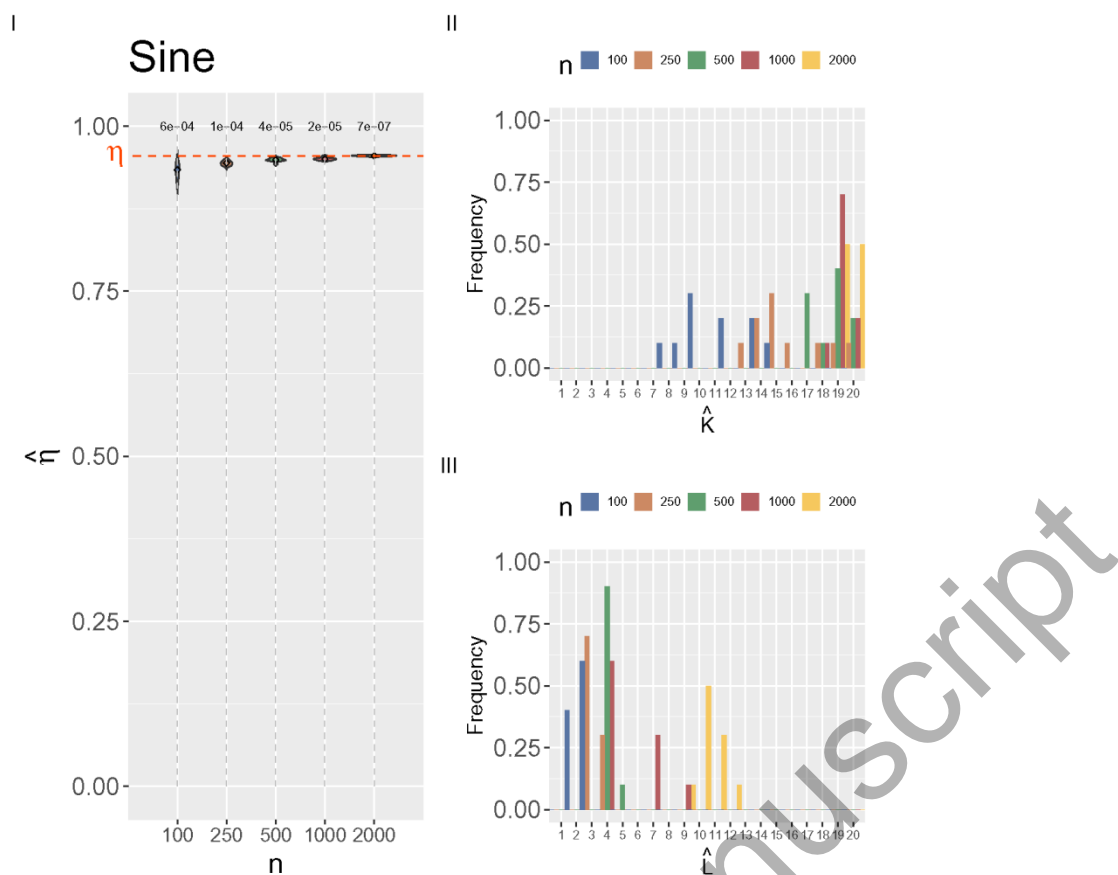


Fig. D.11 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

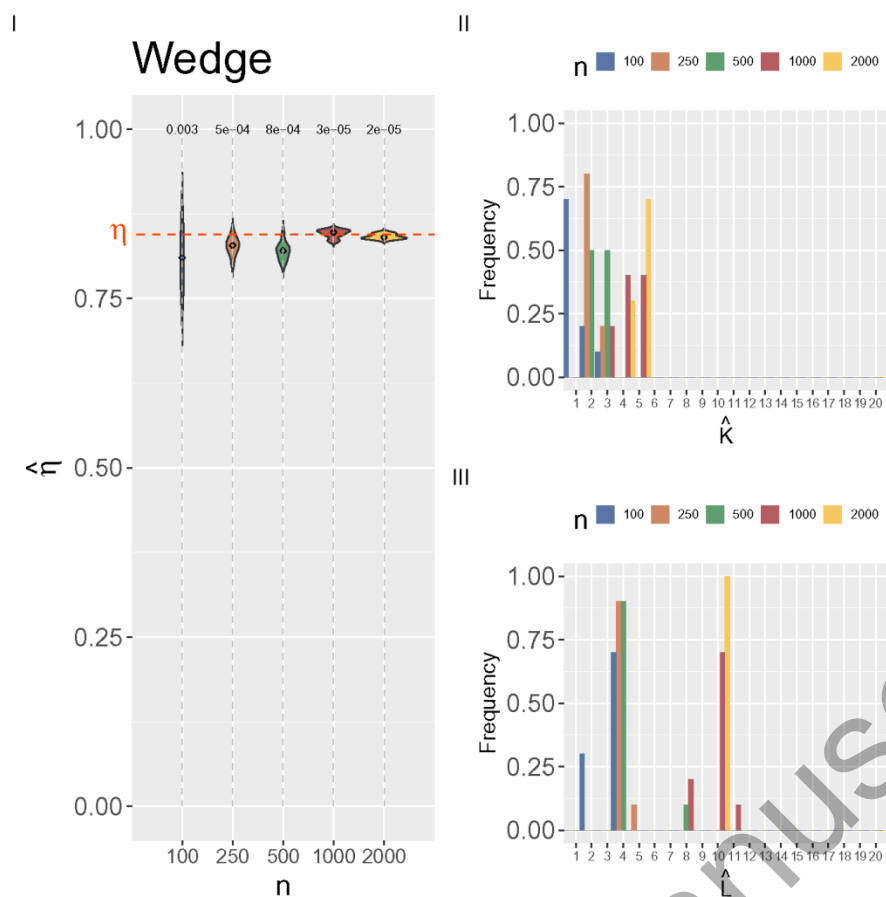
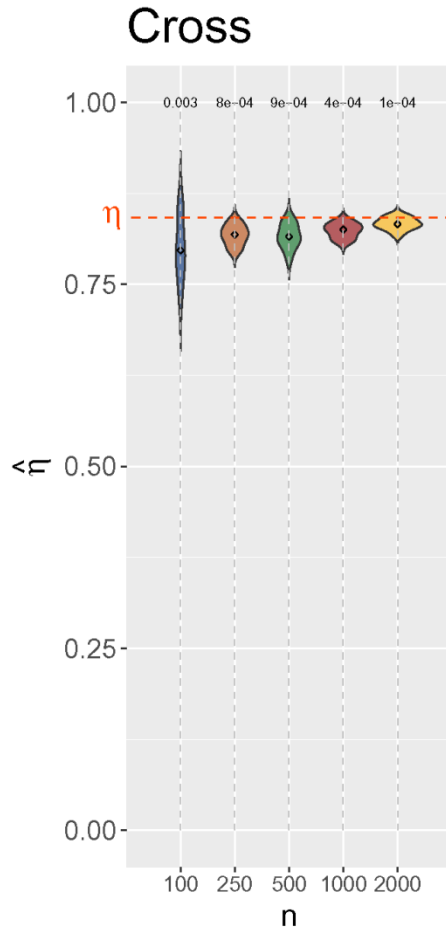



Fig. D.12 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III






Fig. D.13 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III






Fig. D.14 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III






Fig. D.15 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

I



II



III




Fig. D.16 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

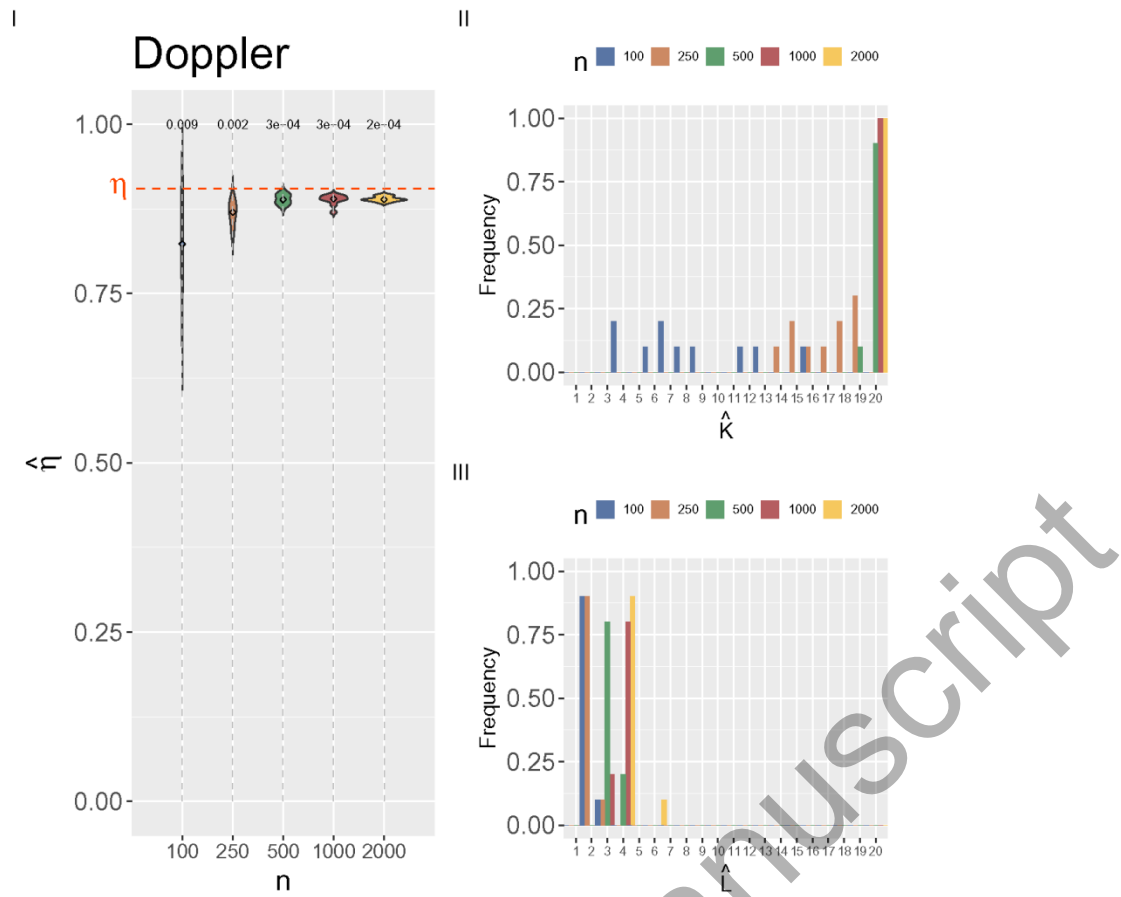


Fig. D.17 $M = 1,000$; $n_{\infty} = 20,000$; $\alpha = 0.7$, $\beta = 0.01$

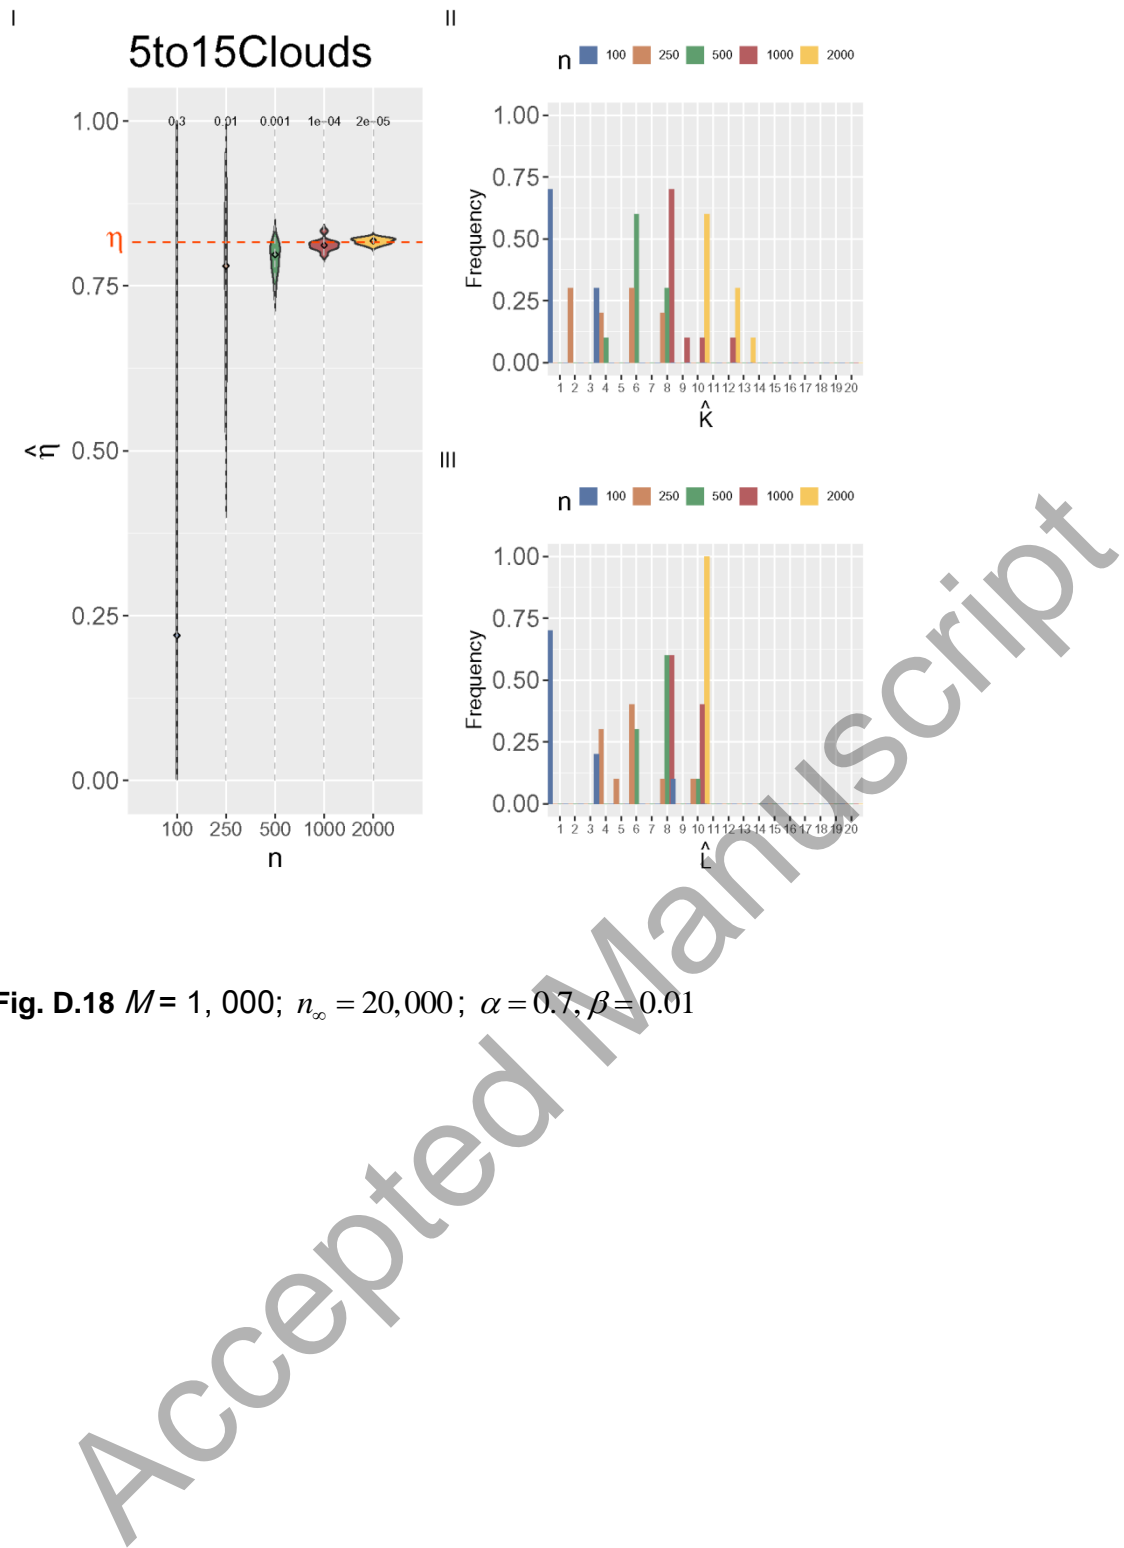


Table 5.2 Empirical power (in %) of 7 tests of independence for the Peano ($d=4$) structure of pure dependence, for increasing sample size n . (The entry labelled ‘> time’ was not computed due to a prohibitive computing time for a sample size of 500.) Number of Monte-Carlo replications: $M=10,000$.

n	dCov	HHG	\mathcal{G}^2	RCD	GMR	$\hat{\eta}_0$	$\hat{\eta}$
100	5.14	5.28	2.77	5.46	4.88	7.34	5.30
200	5.75	5.44	1.19	4.95	1.46	14.86	6.81
300	5.00	5.13	0.78	5.38	0.026	27.79	7.81
400	5.38	5.31	0.58	5.45	$6e-04$	58.12	8.86
500	4.83	4.88	0.67	5.61	> time	86.06	8.72

Table 5.3 Empirical power (in %) of 7 tests of independence for the 'bisection expanding cross' ($d = 3$) structure of pure dependence, for increasing sample size n . Number of Monte-Carlo replications: $M = 10,000$.

n	dCov	HHG	\mathcal{G}^2	RCD	GMR	$\hat{\eta}_0$	$\hat{\eta}$
50	4.94	8.97	4.03	7.33	5.62	28.83	5.71
100	4.68	12.44	2.43	15.55	28.17	96.06	8.46
150	4.60	15.82	1.33	44.08	94.46	100.00	10.79
200	5.16	20.68	1.00	89.84	100.00	100.00	17.92
300	5.25	34.50	0.71	100.00	100.00	100.00	67.54

Accepted Manuscript

Table 6.1 Density of seabirds and fish (individuals per hectare of island, data truncated to the nearest integer) around $n = 12$ islands of the Chagos Archipelago. (Untruncated data are provided in package `HellCor` via the command `data(Chagos)`).

[illegible]

Table D.1 Fifteen scenarii of dependence. Functions Ψ_1 and Ψ_2 from (D.1):

$\Psi_1 = F_Y^{-1}$ and $\Psi_2 = F_Z^{-1}$, where $Y \stackrel{d}{=} \psi_1(U_1, U_2)$, $Z \stackrel{d}{=} \psi_2(U_1, U_2)$ and where the random variables U_1 and the univariate components of the ρ -dimensional random vector U_2 are mutually independent uniform random variables taking values in $[0, 1]$. A bit of jittering is added to the ‘4 Clouds’ and ‘5 to 15 Clouds’ scenarios to ensure that the random variables X_1 and X_2 have a continuous part so that our measure is properly defined.

Name	ρ	$\psi_1(u_1, u_2)$	$\psi_2(u_1, u_2)$
W	0	$(2u_1 - 1)$	$4 \left\{ \left[(2u_1 - 1)^2 - 0.5 \right]^2 \right\}$
Diamond	1	$2u_1$	$-1 + (2I_{\{u_2 > 0.5\}} - 1)(1 - 2u_1 - 1)$
Parabola	0	$(2u_1 - 1)$	$(2u_1 - 1)^2 / 2$
2Parabolas	1	$(2u_1 - 1)$	$(2u_1 - 1)^2 (2I_{\{u_2 > 0.5\}} - 1)$
Circle	0	$\sin(2\pi u_1 - \pi)$	$\cos(2\pi u_1 - \pi)$
4Clouds	0	$(2I_{\{u_1 > 0.5\}} - 1) + u_1 / 100$	$(2I_{\{u_1 > 0.5\}} - 1) + u_1 / 100$
Cubic	0	$(2u_1 - 1)^3$	$4(2u_1 - 1)^3 + (2u_1 - 1)^2 - 4(2u_1 - 1)$
Sine	0	$4u_1$	$4\sin(4\pi u_1)$
Wedge	1	u_1	$u_1(2I_{\{u_2 > 0.5\}} - 1)$
Cross	1	u_1	$(2u_1 - 1)(2I_{\{u_2 > 0.5\}} - 1)$
Spiral	0	$6\pi u_1 \cos(6\pi u_1) / 3$	$6\pi u_1 \sin(6\pi u_1) / 3$
4Circles	2	$\sqrt{2} \cos(2\pi u_1) + (2I_{\{u_2^{(1)} > 0.5\}} - 1)$	$\sqrt{2} \sin(2\pi u_1) + (2I_{\{u_2^{(2)} > 0.5\}} - 1)$
Heavisine	0	$4u_1$	$y / \text{var}(y)$ with
			$y = 4\sin(4\pi u_1) - \text{sign}(u_1 - 0.3) - \text{sign}(0.72 - u_1)$
Doppler	0	$10u_1$	$y / \sqrt{\text{var}(y)}$ with $\delta = 0.05$ and
			$y = \sqrt{u_1(1 - u_1)} \sin \{ [2\pi(1 + \delta)] / (u_1 + \delta) \}$
5to15Cloud	0	$u_1 / 100 - (I_1 + I_2 - I_4 - I_5)$	$u_1 / 100 - (I_1 - I_2 - I_4 + I_5)$

Name	ρ	$\psi_1(u_1, u_2)$	$\psi_2(u_1, u_2)$
ds			
		with $I_j = I_{\{(j-1)/5 \leq u_1 < (j/5)\}}, j = 1, \dots, 5$	with $I_j = 1_{\{(j-1)/5 \leq u_1 < (j/5)\}}, j = 1, \dots, 5$

Accepted Manuscript

Table D.2 Power (in %) with $n = 20$, $M = 10,000$, $\alpha = 0.7$, $\beta = 0.01$. (Processing time: 46h.)

	dCov	HHG	\mathcal{G}^2	RCD	GMR	$\hat{\eta}_0$	$\hat{\eta}$
W	29.8	69.3	40.2	81.0	31.8	59.7	63.7
Diamond	2.9	25.4	22.8	38.5	13.2	27.7	19.4
Parabola	61.6	97.9	32.5	96.5	56.6	83.5	96.8
2Parabolas	21.1	79.0	7.4	62.7	18.4	44.9	39.5
Circle	3.8	13.8	15.8	28.7	1.8	24.0	13.3
4Clouds	99.9	99.9	77.5	99.5	0.0	73.3	99.2
Cubic	81.5	93.8	14.0	98.6	4.3	83.5	97.7
Sine	38.7	54.2	25.6	66.8	35.4	53.4	70.5
Wedge	38.6	66.3	14.4	61.0	14.0	37.7	61.2
Cross	13.2	53.2	7.0	50.3	9.9	32.7	25.3
Spiral	6.1	6.6	5.5	6.7	5.0	7.3	5.8
4Circles	4.7	4.6	5.1	5.4	4.3	5.3	4.9
Heavisine	48.6	67.7	32.0	79.6	61.9	60.0	74.0
Doppler	21.3	30.1	29.3	33.8	14.6	25.9	21.8
5to15Clouds	7.4	21.8	9.7	30.3	2.1	20.6	9.3

Table D.3 Power (in %) with $n = 50$, $M = 10,000$, $\alpha = 0.7$, $\beta = 0.01$. (Processing time: 57h.)

	dCov	HHG	\mathcal{G}	RCD	GMR	$\hat{\eta}_0$	$\hat{\eta}$
W	94.5	100.0	78.6	100.0	96.0	99.9	100.0
Diamond	9.6	93.5	81.1	96.8	53.5	70.6	95.5
Parabola	99.8	100.0	83.2	100.0	99.8	99.9	100.0
2Parabolas	50.5	100.0	1.8	99.6	81.6	91.3	99.0
Circle	7.1	72.0	61.1	88.3	1.5	68.0	88.2
4Clouds	100.0	100.0	100.0	100.0	0.0	99.5	100.0
Cubic	100.0	100.0	58.8	100.0	2.8	99.9	100.0
Sine	93.8	100.0	82.5	100.0	97.4	99.7	100.0
Wedge	95.3	99.8	52.3	99.2	64.4	79.3	99.8
Cross	25.7	98.2	1.9	98.4	51.4	73.5	88.9
Spiral	8.3	10.0	6.6	11.4	2.5	12.0	6.7
4Circles	4.9	4.9	6.6	5.4	4.1	7.6	4.8
Heavisine	98.8	100.0	84.5	100.0	100.0	99.8	100.0
Doppler	65.9	94.4	93.1	91.4	89.2	71.5	70.9
5to15Clouds	11.1	55.1	2.9	86.7	0.0	58.3	13.8

Notes

¹An earlier version of this example was used in Boglioni-Beaulieu (2016).

²Source: <http://www.people.fas.harvard.edu/~junliu/Gs/Gs.html>