

## Nonparametric Testing for Heterogeneous Correlation

Stephen Bamattre<sup>1</sup>, Rex Hu<sup>1</sup> and Joseph S. Verducci<sup>1</sup>  
The Ohio State University, 1958 Neil Ave, Columbus OH 43210

### Abstract

In the presence of weak overall correlation, it may be useful to investigate if the correlation is significantly and substantially more pronounced over a subpopulation. Two different testing procedures are compared. Both are based on the rankings of the values of two variables from a data set with a large number  $n$  of observations. The first maintains its level against Gaussian copulas; the second adapts to general alternatives in the sense that the number of parameters used in the test grows with  $n$ . An example illustrates how the test statistics may be used to discover pairs of genes whose expressions in rat hepatocyte cells are substantially more correlated over those cells that have been exposed to certain classes of chemicals.

**Keywords:** Absolute rank differences; Beta distribution; Frank copula; Gaussian copula; Kendall's tau; Mallows' model; Multistage ranking model; Permutations; Seriation.

### 1. Introduction

The goal of this paper is to offer new methods for discovering association between two variables that is supported only in a subpopulation. For example, while higher counts of HDLs are generally associated with lower risk of myocardial infarction, researchers (Voight et al., 2012; Katz, 2014) have found subpopulations that do not adhere to this trend. In marketing, subpopulations of designated marketing areas (DMAs) in the US respond differentially to TV advertising campaigns, and the identification of DMAs that are sensitive to ad exposure enables efficient spending of ad dollars. In preclinical screening of potential drugs, various subpopulations of chemicals elicit concomitant responses from sets of hepatocyte genes, which can be used to discover gene networks that breakdown classes of drugs, without having to pre-specify how the classes are formed. The new methods thus lead to a whole new approach to analysis of large data sets.

When covariates are available, regression analysis classically attempts to identify a supporting subpopulation via interaction effects, but these may be difficult to interpret properly. In the presence of overall correlation, it may be useful to investigate directly if the correlation is significantly and substantially more pronounced over a subpopulation. This becomes feasible when representatives of supporting subpopulations are embedded in large samples. The novel statistical tests described in this paper are designed to probe large samples to ascertain if there is such a subpopulation.

The general setting is this: A large number  $n$  of observations are sampled from a bivariate continuous distribution. The basic assumption is that the population consists of two subpopulations. In one, the two variables are positively (or negatively) associated; in the other, the two variables are independent. While some distributional assumptions are required even to define the notion of homogeneous association, the underlying intent is to make the tests robust to assumptions about the distributions governing both the null and alternative hypotheses.

Notation for the rest of the paper is as follows: Let  $X \sim F$  and  $Y \sim G$  have joint, continuous distribution  $H$ . For any sample  $\{(x_i, y_i) \mid i = 1, \dots, n\}$ , the empirical marginal distributions are defined by

$$\hat{F}_n(x) = \frac{1}{n} \sum_{i=1}^n 1\{x_i \leq x\} \quad \text{and} \quad \hat{G}_n(y) = \frac{1}{n} \sum_{i=1}^n 1\{y_i \leq y\}.$$

The ranking  $\pi$  of the sample  $\{x_i \mid i = 1, \dots, n\}$  is the function  $\pi: \{x_i \mid i = 1, \dots, n\} \rightarrow \{1, \dots, n\}$  defined by

$$\pi(x_i) = \sum_{j=1}^n 1\{x_i \leq x_j\}.$$

The corresponding ranking of  $\{y_i \mid i = 1, \dots, n\}$  is denoted by  $\nu$ . Spearman's footrule distance with a sample  $\{(x_i, y_i) \mid i = 1, \dots, n\}$  is defined through the sample rankings as

$$d_S = \sum_{i=1}^n |\pi(x_i) - \nu(y_i)|.$$

The Kendall distance associated with the sample is defined as

$$D_K([x_1, \dots, x_n], [y_1, \dots, y_n]) = \sum_{i < j} 1\{(x_i - x_j)(y_i - y_j) < 0\} = \sum_{i < j} 1\{(\pi(x_i) - \pi(x_j))(\nu(y_i) - \nu(y_j)) < 0\} = d_K(\pi, \nu),$$

which also depends only on the rankings  $\pi$  and  $\nu$  of the sample  $\{x_i\}$  and  $\{y_i\}$ . Mallows (1957) model for rankings takes the form

$$P_\phi(\nu \mid \pi) = C(\phi) e^{-\phi d_K(\pi, \nu)}$$

where the normalizing constant  $C(\phi)$  has a tractable form (Fligner and Verducci, 1986) known as a Poincaré polynomial (Diaconis and Graham, 2000). Distributional forms for the data are in terms of copulas:

$$C_H(F(X), G(Y)) = H(X, Y),$$

which are distribution functions on the unit square, having uniform margins. Two copulas play a fundamental role in motivating the tests: the Gaussian copula and the Frank Copula. If  $(X, Y)$  has a bivariate normal distribution  $H$  with correlation  $\rho$ , then its corresponding copula is

$$C_\rho(u, v) = \Phi_2(\Phi^{-1}(u), \Phi^{-1}(v); \rho)$$

where  $\Phi$  is the standard normal CDF. The bivariate distributions  $C_\rho$  and  $\Phi_2$  are indexed solely by the underlying correlation  $\rho$ . The Frank copula (Frank, 1979; Genest, 1987) has the form

$$C_\theta(u, v) = -\frac{1}{\theta} \log \left( 1 + \frac{(e^{-\theta u} - 1)(e^{-\theta v} - 1)}{(e^{-\theta} - 1)} \right).$$

The next two sections describe two new tests for detecting subpopulations that support association: the Components of Spearman's Footrule (CSF) test and the Components of Kendall's Tau (CKT) test. The CSF test is scaled according to a Gaussian copula and the CKT test is scaled according to a Frank copula. The CSF test is computationally fast, and the CKT test adapts to a large variety of alternatives. The following two sections cover their performance under simulations. Concluding remarks are in last section.

## 2. Components of Spearman's Footrule (CSF)

While Spearman's footrule (Diaconis and Graham, 1977) measures the overall disarray in a sample, the distribution of individual absolute rank differences

$$d_i = |\pi(x_i) - \nu(y_i)|$$

proves to be very useful in detecting subsamples with distinctly less disarray than would be expected under homogeneous association. Because the rankings depend on the whole sample, the  $\{d_i\}$  are not independent. Nevertheless, we loosely define their *empirical distribution* as

$$S_n(d) = \frac{1}{n} \sum_{i=1}^n 1\{d_i \leq d\}.$$

As a step toward determining asymptotic forms for this distribution, we offer the following lemmas:

**Lemma 1.** For any sample  $\{(X_i, Y_i) \mid i = 1, \dots, n\}$  from a joint distribution  $H$ , let  $(X, Y)$  be a newly sampled observation. Then, for rankings  $\pi$  and  $\nu$  for the extended sample, the asymptotic distribution of

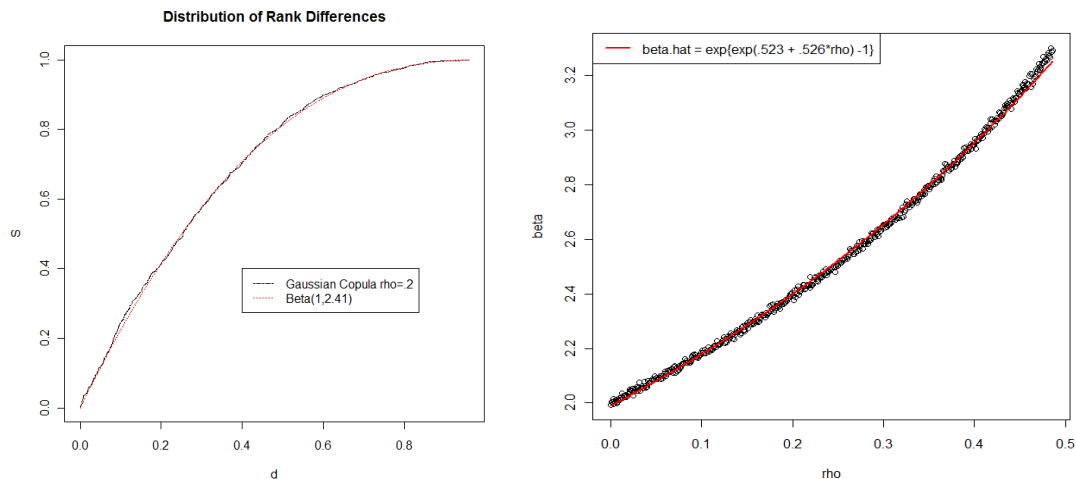
$$\left[ \frac{\pi(X)}{n+1}, \frac{\nu(Y)}{n+1} \right] \approx \left[ \hat{F}(X), \hat{G}(Y) \right]$$

is the underlying copula  $C_H[F(X), G(Y)]$  of  $H$ .

**Lemma 2.** The asymptotic distribution of rank differences under independence is Beta(1,2).

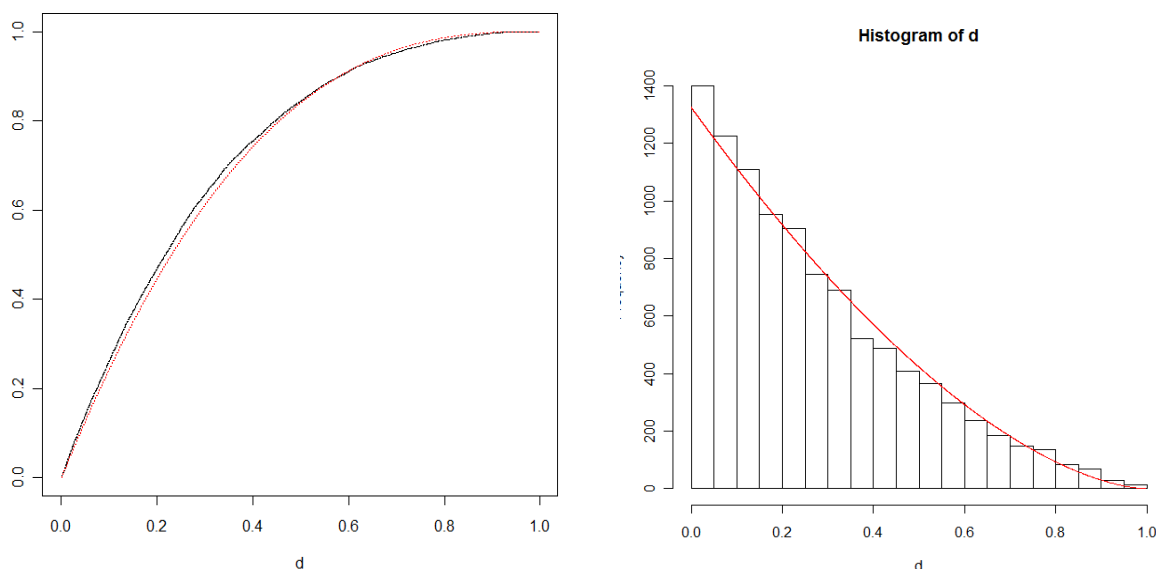
These lemmas lead to the following

**Proposition 1.** Under a Gaussian( $\rho$ ) copula,  $S_n$  converges to a Beta(1,  $\beta(\rho)$ ) distribution. Simulations with  $n = 1000$  affirm the proposition and produce a smooth curve for  $\beta(\rho)$ . See Figure 1.



**Figure 1:** Left panel shows the close agreement of  $S_n$  with Beta(1,  $\beta(\rho)$ ) when sampling  $n = 1000$  observations from a Gaussian( $\rho = .2$ ) copula. The right panel illustrates the  $\beta(\rho)$  curve for  $0 < \rho < .5$ .

The null hypothesis is that  $(X,Y)$  have a Gaussian copula. The alternative is that  $(X,Y)$  come from a mixture of two subpopulations in which under one they are independent, and under the other they are positively associated. To test for negative association, simply replace  $Y$  by  $-Y$ . No particular form is assumed for the positively associated subpopulation, but it is informative to examine the case where this component is Gaussian. Figure 2 illustrates  $S_n$  and its histogram under such a mixture.



**Figure 2:** (Standardized) Distribution of Rank Differences Under Mixture of Gaussian (.6) and Independent Copulas with overall correlation .3 compared to Beta(1,2.65)

Because the differences in distributions under the null and alternative are small, large samples are required to distinguish the two. As noted from the histogram in Figure 2, most of the distinguishing information is contained at the low end of the distribution. *This makes sense because a subpopulation supporting positive association should have a surplus of points where the ranks of  $X$  and  $Y$  closely agree.* Thus a test statistic based absolute ranked differences should emphasize the lower order statistics. Such statistics come under the heading of L-statistics. It is possible to tailor a test toward alternative features of interest such as proportionate size of the subpopulation and the strength of association within it. Exact distributions of partial or weighted sums of absolute rank differences are quite complicated due to dependencies (Sen, et al. 2011), even under the null hypothesis of independence. A very simple general purpose test statistic is

$$T_S = \sum_{i=1}^n 1\{\left|\pi(x_i) - \nu(y_i)\right| < .2\}.$$

Using the observed overall correlation  $r$  in place of  $\rho$ , the null distribution of  $T_S$  may be simulated under the Gaussian copula or approximated as a Binomial test statistic using the probability from the Beta(1, $\beta$ ) as in Proposition 1. In the later case, ignoring weak dependencies, the .05 level test has power of 80% of detecting a Gaussian subpopulation of 25% with  $r = .8$  for  $n = 1000$ .

### 3. Components of Kendall's Tau (CKT)

Although the CSF test is both simple and computationally efficient, it has a conceptual shortcoming arising from the use of Spearman's footrule distance to characterize association in a subpopulation. The issue is that the components of the footrule distance in the subpopulation depend on the encompassing population; that is, when the sample is a full population, with associated subpopulation  $\Omega$ , the component set from the footrule from  $\Omega$

$$\{d_i; i \in \Omega\} = \{|\pi(x_i) - \nu(y_i)|; i \in \Omega\}$$

depends heavily on the rankings  $\pi$  and  $\nu$  determined by the full population. In contrast, the component set from Kendall's distance depends only on the relative rankings within  $\Omega$ , which may be constructed from just on the original values in  $\Omega$ . That is,

$$\{1[(\pi(x_i) - \pi(x_j))(\nu(y_i) - \nu(y_j)) < 0]; i, j \in \Omega\} = \{1[(x_i - x_j)(y_i - y_j) < 0]; i, j \in \Omega\}.$$

Thus the subpopulation discordances (components of Kendall's distance) do not depend upon the embedding population, whereas the subpopulation disarray (components of Spearman's footrule distance) do. This invariance has a number of beneficial properties, such as allowing the CKT test to retain power in situations where the ranges of the  $\{X_i\}$  and  $\{Y_i\}$  values in the subpopulation are more restricted than those in the full population.

The notion of homogeneous association based on Kendall's distance differs from that based the Spearman's footrule used for the CSF test. In this case the natural null hypothesis should be a distribution depending only on Kendall's distance. Furthermore it should have the greatest entropy for a given value of Kendall's tau because this formulation would attribute as much variability as possible to the null distribution, making it a conservative (least favorable) test (Lehmann and Romano, 2006). To construct a distribution has this structure, simply sample from an arbitrary copula, and then reorder the Y-values according to a permutation  $\nu(\mathbf{Y})$  sampled independently from a Mallows model centered at the ranking  $\pi(\mathbf{X})$  of the X values. Quite remarkably, any such process asymptotically leads to a Frank copula. Proposition 2, based on Starr (2009), gives a precise statement.

**Proposition 2.** Let  $\{(X_i, Y_i) \mid i=1, \dots, n\}$  be independent samples from a distribution  $H$  with continuous marginals  $F$  and  $G$ , and associated copula  $C$  with continuous partial derivatives. Let  $\pi(\mathbf{X})$  be the ranking of  $\mathbf{X}=[X_1, \dots, X_n]$  and  $\nu(\mathbf{Y})$  be the ranking of  $\mathbf{Y}=[Y_1, \dots, Y_n]$ . Assume that for all  $n$  sufficiently large, the conditional distribution of  $\nu(\mathbf{Y})$  given  $\pi(\mathbf{X})$  is Mallows, with center at  $\pi(\mathbf{X})$  and scale  $\phi_n$ . If  $\phi_n \rightarrow 0$ , and there exists  $\theta \neq 0$  such that

$$n(1 - e^{-\phi_n}) \rightarrow \theta,$$

then  $C$  is the Frank copula  $C_\theta$ .

**Proof.** First, it is necessary to establish radial symmetry of a copula that has  $\nu(\mathbf{Y})$  given  $\pi(\mathbf{X})$  being a Mallows distribution. The *pseudo-observations* for each pair  $(X_i, Y_i)$  are defined as functions of the pair and the empirical margins:

$$(\hat{U}_i, \hat{V}_i) = \frac{n}{n+1} (\hat{F}_n(X_i), \hat{G}_n(Y_i)).$$

These are functions of the rankings  $\pi(\mathbf{X})$  and  $\nu(\mathbf{Y})$ :

$$\hat{U}_i = 1 - \frac{\pi(X_i)}{n+1} \quad \hat{V}_i = 1 - \frac{v(Y_i)}{n+1}$$

By the symmetry of the Mallows model, the joint distribution of the pseudo-observations  $(\hat{U}_1, \hat{V}_1), \dots, (\hat{U}_n, \hat{V}_n)$  is identical to the joint distribution of  $(1 - \hat{U}_1, 1 - \hat{V}_1), \dots, (1 - \hat{U}_n, 1 - \hat{V}_n)$ . Consider empirical distributions based on these observations (Genest and Nešlehová, 2014):

$$\hat{C}_n(u, v) = \frac{1}{n} \sum_{i=1}^n I[\hat{U}_i \leq u, \hat{V}_i \leq v]$$

$$\hat{D}_n(u, v) = \frac{1}{n} \sum_{i=1}^n I[1 - \hat{U}_i \leq u, 1 - \hat{V}_i \leq v]$$

Since  $H$  has continuous marginals and  $C$  has continuous partial derivatives, then Fermanian *et al.* (2004) established that  $\hat{C}_n$  is a consistent estimator of the copula  $C$ , and likewise  $\hat{D}_n$  is a consistent estimator of the *survival copula*  $\bar{C}$ :

$$\bar{C}(u, v) = u + v - 1 + C(1 - u, 1 - v)$$

Hence,  $C = \bar{C}$ , which implies that the copula  $C$  is radially symmetric (Nelsen 2006, pg. 37). Now, since  $C$  is radially symmetric, an asymptotically equivalent definition of the *empirical copula* is:

$$\check{C}_n(u, v) = \frac{1}{n} \sum_{i=1}^n I \left[ \frac{\pi(X_i)}{n} \leq u, \frac{v(Y_i)}{n} \leq v \right] = \frac{1}{n} \sum_{i=1}^n \delta_{(\pi(X_i)/n, v(Y_i)/n)}$$

Which places mass of  $\frac{1}{n}$  on each random point  $\left(\frac{\pi(X_i)}{n}, \frac{v(Y_i)}{n}\right) \in [0,1]^2$ . This empirical copula is expressed by a point process (Starr, 2009): For  $n \in \mathbb{N}$ ,

$$\mu_n(B, \omega) = \frac{1}{n} \sum_{i=1}^n I \left[ \left( \frac{\pi(X_i(\omega))}{n}, \frac{v(Y_i(\omega))}{n} \right) \in B \right]$$

for each bounded Borel set  $B \subseteq \mathbb{R}^2$ .

By assumption, the regularity conditions on the Mallows scale are satisfied as  $n \rightarrow \infty$ :

$$\phi_n \rightarrow 0, \quad \exists \theta \in \mathbb{R} \setminus \{0\} \ni n(1 - e^{-\phi_n}) \rightarrow \theta$$

Under these conditions, the primary result of Starr (2009) is applied: As  $n \rightarrow \infty$ , the random measures  $\mu_n(\cdot, \omega)$  weakly converge to the measure  $\mu_\theta$ , defined by:

$$d\mu_\theta(u, v) = \frac{(\theta/2) \sinh(\theta/2)}{(e^{\theta/4} \cosh(\theta[u-v]/2) - e^{-\theta/4} \cosh(\theta[u+v-1]/2))^2} I_{[0,1]^2}(u, v) \partial u \partial v$$

Simply converting the trigonometric functions to exponential form and simplifying yields:

$$d\mu_\theta(u, v) = \frac{\theta(1 - e^{-\theta})e^{-\theta(u+v)}}{(1 - e^{-\theta} - (1 - e^{-\theta u})(1 - e^{-\theta v}))^2} I_{[0,1]^2}(u, v) \partial u \partial v$$

By recognition, the limiting measure is that of the (Frank) Copula  $C_\theta$ :

$$d\mu_\theta(u, v) = c_\theta(u, v) I_{[0,1]^2}(u, v) \partial u \partial v = \left[ \frac{\partial^2}{\partial u \partial v} C_\theta(u, v) \right] I_{[0,1]^2}(u, v) \partial u \partial v$$

Recall,  $\tilde{C}_n$  is a consistent estimator of the underlying copula  $C$ , and converges weakly to  $C_\theta$ , so we conclude that  $C = C_\theta$ .  $\triangle$

Pursuing this result further allows for inspection of the adequacy of the asymptotic result for finite samples. A function  $\phi(\theta)$  for matching the Mallows  $\phi$  parameter to the Frank( $\theta$ ) may be obtained by equating expressions for  $\tau_\phi$  and  $\tau_\theta$  from these models. For any Archimedean copula, there is a relatively simple formula  $\tau = 4E[C(U, V)] - 1$  (MacKay and Genest, 1986); for the Frank copula, a specialized form (Nelsen, 2006, p. 171; Genest, 1987) is

$$\tau_\theta = 1 - \frac{4}{\theta} [1 - D(\theta)] = 1 - \frac{4}{\theta} \left[ 1 - \frac{1}{\theta} \int_0^\theta \frac{t}{e^t - 1} dt \right]$$

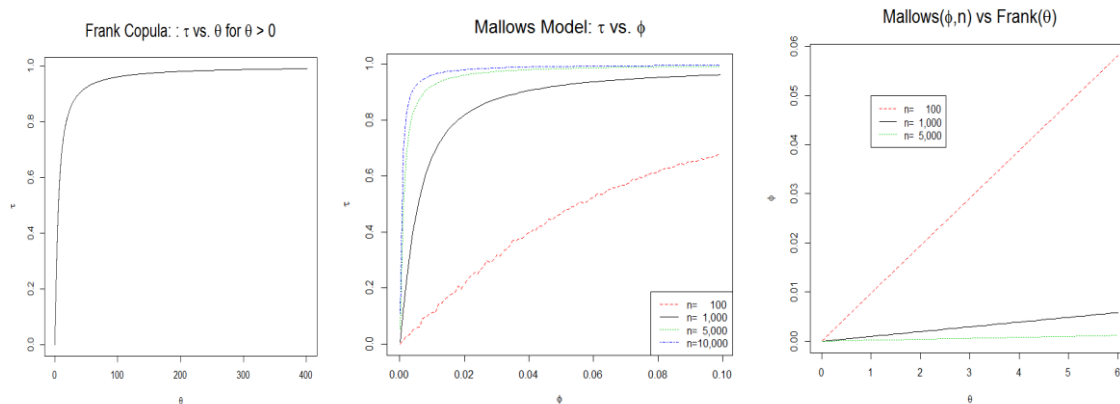
where the scaled integral  $D(\gamma)$  is known as the Debye-1 function, available in the “gsl” (Gnu Scientific Library) package of R. For the Mallows model,

$$\tau_\phi = \frac{2}{\pi} \tan^{-1}(.18n\phi)$$

Equating  $\tau_\theta$  and  $\tau_\phi$  leads to the relationship

$$\phi = \frac{100}{18n} \tan \left[ \frac{\pi}{2} \left\{ 1 - \frac{4}{\theta} [1 - D(\theta)] \right\} \right] \approx \frac{.9694}{n} \theta.$$


Empirical evidence for the applicability of Proposition 2 comes in two stages: 1) The distribution of Kendall’s Distance under Frank( $\theta$ ) and under Mallow( $\phi(\theta)$ ) both converge to the same normal distribution; 2) As  $n$  gets large the product density of the sample under Frank( $\theta$ ) converges to an increasing function of the Kendall’s Distance between  $\pi(\mathbf{X})$  and  $\nu(\mathbf{Y})$  of the sample.



**Figure 3.** Scale relationships. Left: Kendall's  $\tau$  vs. the Frank scale  $\theta$ , center: Kendall's  $\tau$  vs. the Mallows scale  $\phi$  for  $n=100, 1000, 10000$ , right:  $\theta$  vs.  $\phi_n$  for each  $n$ .

Figure 4 illustrates results from the following confirmatory experiment:

- Generate 200 sets of 1000 points from a Frank copula
- Compute the Kendall distance  $D$  and the Frank density  $d$  for each set
- Plot  $d$  vs  $D$  on a log-log scale



**Figure 4.** Association between  $X$  and  $Y$  in a Frank copula approaches a Mallows model for the ranking  $\nu(\mathbf{Y})$  of a large sample centered at  $\pi(\mathbf{X})$ . Left panel: approximate normality of Kendall's Distance  $D$  under a Frank copula; Center panel: agreement of asymptotic normal forms of  $D$  under Frank and Mallows models; Right panel: log-log plot of Kendall's Distance vs. Frank density.

Note also that the Frank copula is radially symmetric [ $C(u, v) = u+v-I+ C(1-u, 1-v)$ ], which is a necessary condition for the density of a sample to depend only on its Kendall distance.

With the assurance that there are copulas with the conditional distribution of  $\nu(\mathbf{Y})$  given  $\pi(\mathbf{X})$  well approximated by a Mallows model, this becomes the null hypothesis:



$\mathbf{H}_0: \nu(\mathbf{Y}) \circ \pi^{-1}(\mathbf{X}) \sim \text{Mallows}(\theta)$  for some  $\theta > 0$ .

The general alternative against which we would like a test to be sensitive is that there is a subpopulation with high association with the remainder having (little or) no association. The test for heterogeneity should maintain power over a wide variety of alternative distributions for the subpopulation supporting strong association. With these considerations, the alternative hypothesis is formulated as

$\mathbf{H}_A: (F(X), G(Y)) \sim M$  where  $M$  is a mixture of two distributions:  $H_1$  on which  $(F(X), G(Y))$  are independent and  $H_2$  on which  $(F(X), G(Y))$  have  $\tau > 0$ .

To test against such a general alternative, an adaptive model encompassing the Mallows model is adopted, with the number of free parameters in the model increasing with sample size. This components of Kendall's tau (CKT) test proceeds in four steps

1. Fit a Mallows model centered at  $\pi(\mathbf{X})$  to  $\nu(\mathbf{Y})$  and compute the likelihood.
2. Reorder the data points  $\{(X_i, Y_i) \mid i=1, \dots, n\}$  so that Kendall's tau coefficient is decreasing. See Yu *et al.* (2011). Call the reordering  $\sigma$ .
3. Smoothly fit a multistage ranking model to the relative rankings of  $[Y_{\sigma(1)}, \dots, Y_{\sigma(k)}]$  to  $[X_{\sigma(1)}, \dots, X_{\sigma(k)}]$  at each stage  $k$ . See Sampath & Verducci (2013). Compute the likelihood under this (encompassing) model.
4. Use the (Generalized) Likelihood Ratio statistic to test  $H_0$ .

Comments on the four steps:

1. Since Kendall's tau distance is invariant to reordering of observations, this is the same as fitting a Mallows model, centered at ranking  $(\sigma\mathbf{X})$ , to the ranking  $(\sigma\mathbf{Y})$ , where  $\sigma$  is the taupath reordering.
2. The idea of reordering is to put the points displaying the highest amount of association earlier in the sequence in order to identify the subpopulation with highest empirical association. The reordering is not unique. Yu *et al.* (2011) discuss various algorithms.
3. The multistage ranking model decomposes the number of discordances [up to  $(n \text{ choose } 2)$ ] between ranking  $(\sigma\mathbf{Y})$  and ranking  $(\sigma\mathbf{X})$ , as a sum of  $n-1$  variables  $\{V_k\}$  with ranges  $\{0, \dots, k\}$ ,  $k = 1, \dots, n-1$ . The model has likelihood

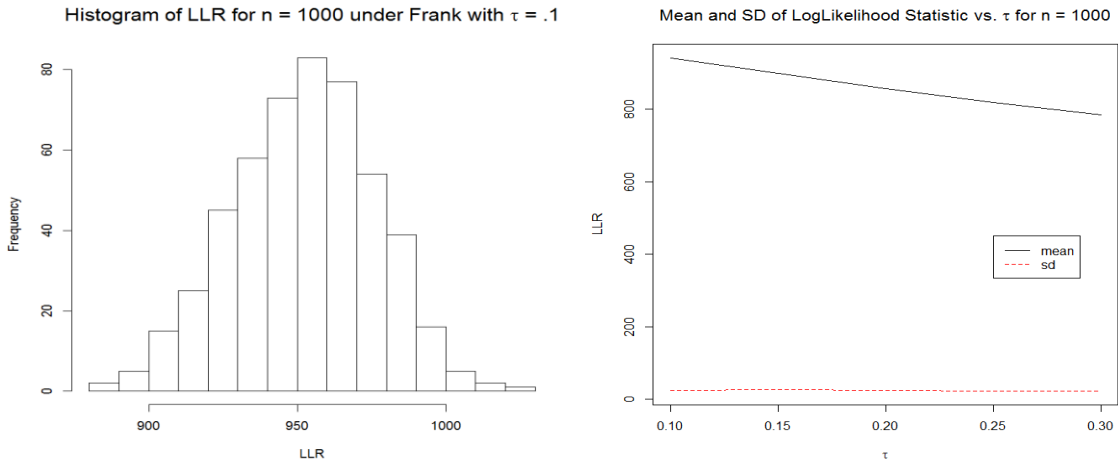
$$L = c(\theta) e^{-\sum \theta_k V_k}$$

which reduces to the likelihood of Mallows model when all component parameters are equal.

4. The conditions needed to justify an asymptotic chi-square distribution for this statistic do not hold in this setting. Currently, we simulate the distribution under the Frank copula to get an appropriate reference. We are working to find a more precise characterization of the LR in this setting.

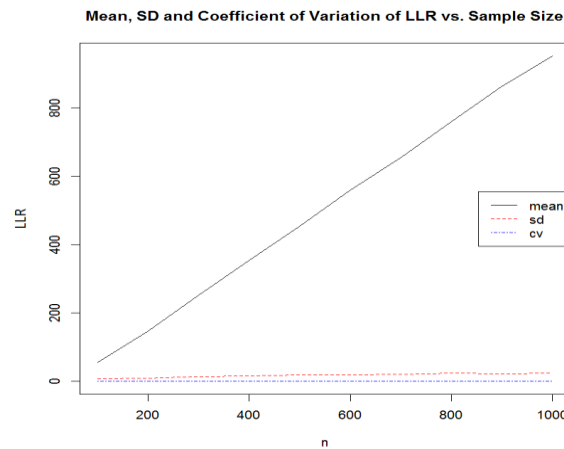
The null distribution of this likelihood ratio appears to be close to normal, with its mean decreasing with the common correlation  $\tau$ , and standard deviation constant. See Figure 5. Note that, for  $n = 1000$ , the

variance of  $2*LLR \approx 2500$  is clearly less than its  $2 * \text{mean}(2*LLR)$  theoretical value for a chi square distribution, which is in the range (3100, 3800) when  $\tau \in (.10, .30)$ .



**Figure 5:** Simulation of Log Likelihood Ratio (LLR) Statistic for CKT test under Frank Copulas. Left Panel: Histogram of 500 simulations of size  $n = 1000$ . Right Panel: decreasing pattern of mean and constancy of standard deviation for LLR under Frank Copulas at different levels of  $\tau$  for  $n = 1000$ .

Instead of fixed  $n$  and varying  $\tau$ , Figure 6 depicts the relationship between  $LLR$  and  $n$  with fixed  $\tau = .1$ .



**Figure 6:** First two moments of LLR are linear in sample size:  $\mu = n - 50$ ;  $\sigma = 7 + .02n$ ;  $cv \rightarrow .02$ .

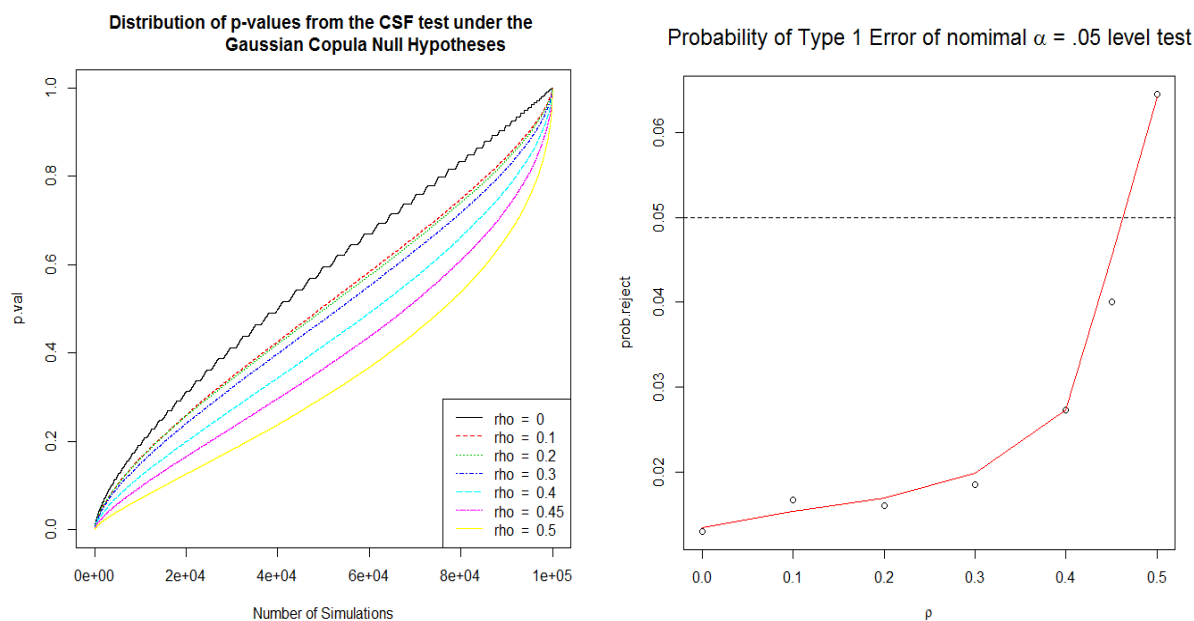
The overall relationship between the moments of LLR and the parameters  $\tau$  and  $n$  is not yet known, but using a practical additive approximation in the range  $.1 < \tau < .30$  and  $500 < n < 3000$ , the basic asymptotic  $\alpha$ -level CKT test has the form: Reject  $H_0$  if

$$Z = \frac{LLR - (n + 20 - 797\hat{\tau})}{.02n + 7} > z_{1-\alpha}$$

where  $\hat{\tau}$  is Kendall's correlation coefficient and  $z_{1-\alpha}$  is the  $(1 - \alpha)^{\text{th}}$  quantile of the standard normal.


#### 4. Simulations for Robustness and Power.

First, performance of the tests is checked by maintenance of levels under various Gaussian and Frank copulas; subsequently power is examined. The CSF test is based on the number of absolute rank differences less than .2. Figure 7 shows the null distributions of p-values for the CSF test applied to samples of size  $n = 1000$  generated 100,000 times under the Gaussian( $\rho$ ) models. These distributions start to become stochastically smaller than uniform for  $\rho > .45$ . Otherwise the test is conservative in the range  $0 < \rho < .45$  and  $0 < \alpha < .05$  as illustrated by the observed number of type 1 errors at the  $\alpha = .05$  level.



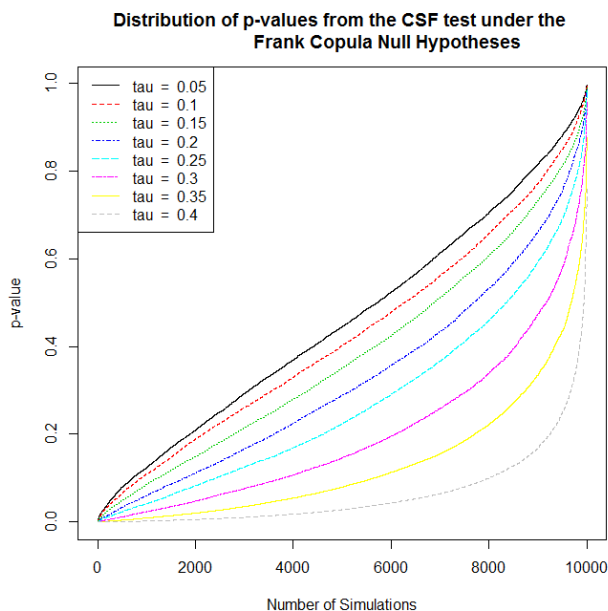
**Figure 7:** CSF results from 100,000 simulation experiments of size  $n = 1000$  for values of  $\rho$  in a Gaussian copula: Left panel—Distribution of p-values for 7 values of the null correlation; Right panel—observed probabilities of Type 1 error for a nominal  $\alpha = .05$  test. The test is conservative for values of  $\rho < .45$ .

Under similar Gaussian copulas, the adjustment of the mean of the log-likelihood for the estimated overall  $\tau$  makes the CKT test behave conservatively for large values of  $\rho$ , but gives highly significant values for  $\tau$  values near 0. See Figure 8, in which, due to computational limitations, lowess-smoothed curves describe the p-distribution based on only 100 simulations. In the presence of very low overall correlation, it is advisable to use the CSF test as a screen for the CKT, which will protect the CKT from finding uneven levels of  $\tau$  association when  $\rho$  association is homogeneous. Again, this tendency toward excess false positives happens only when the overall  $\rho$  association is close to 0. In this case a special test (Sampath and Verducci, 2013) is available for the null hypothesis of independence.



**Figure 8:** CKT results from 200 simulation experiments of size  $n = 1000$  for values of  $\rho$  in a Gaussian copula: Left panel—Distribution of p-values for 7 values of the null correlation; Right panel— observed probabilities of Type 1 error for a nominal  $\alpha = .05$  test. The test is conservative for values of  $\rho > .16$ .

Under a Frank copula, the CSF test behaves properly near independence, but loses its level when  $\tau$  gets large. See Figure 9.

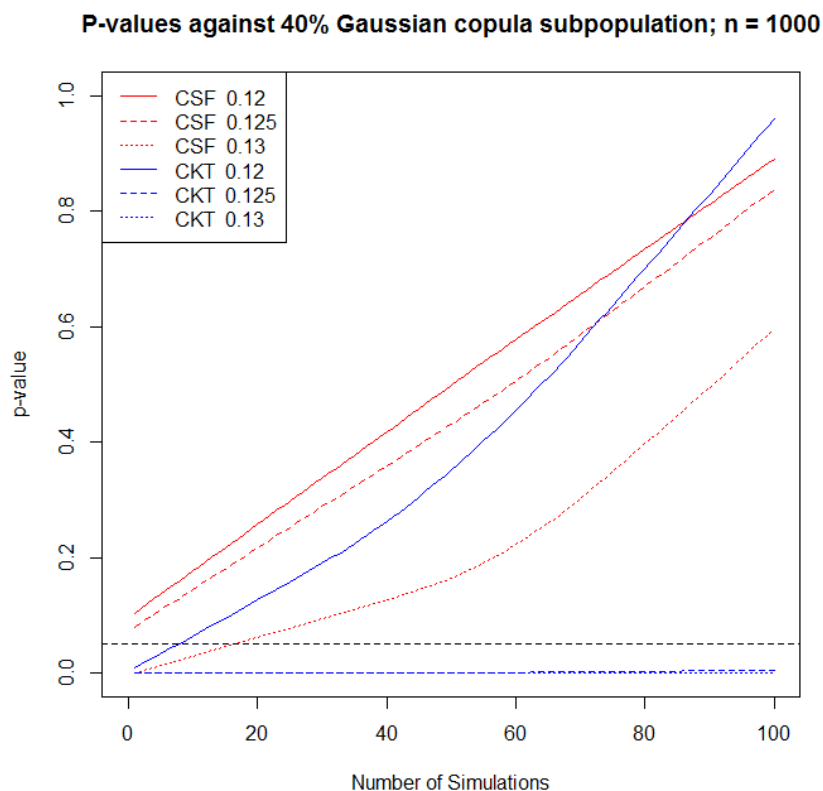


**Figure 9:** Performance of CSF under Frank Copula. As depicted for 10,000 simulations of  $n = 1000$  samples from each Frank copula, the test behaves reasonably for small values of  $\tau$ , but distorts to produce excess false positives when gets  $\tau$  large.

Several factors affect the power curves of both the CSF and CKT tests: *sample size* ( $n$  is fixed at 500 or 1000); *proportionate size of the subpopulation* (fixed at 40%); *strength of association in the subpopulation* ( $\rho, \tau \in \{.7, .8, .9\}$ ); and, most importantly, the *form of the subpopulation*. Against the null hypothesis of a Gaussian copula, the alternative is a mixture of copulas, where the variables are assumed

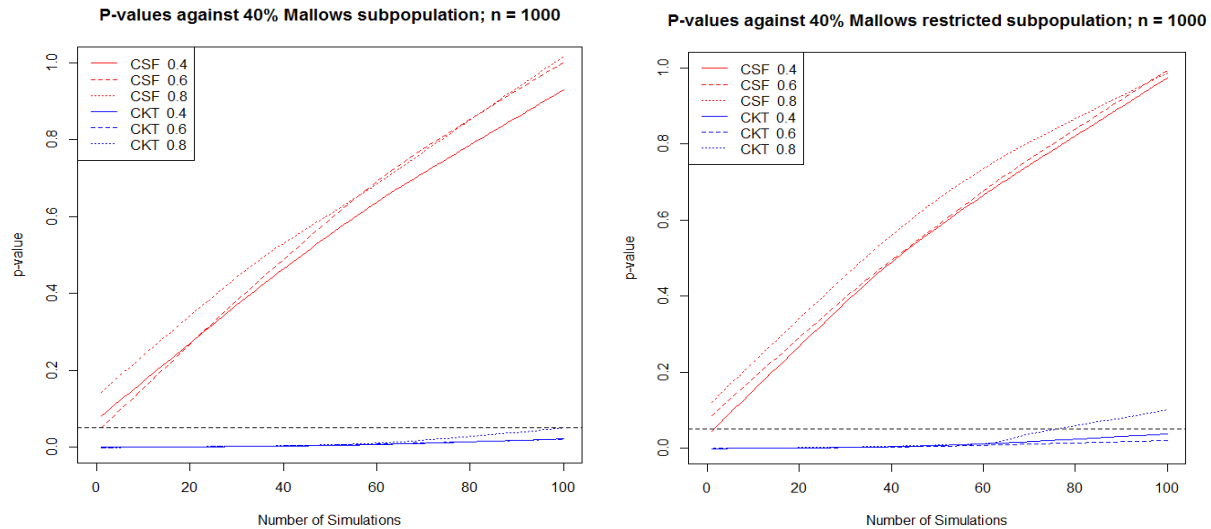
to be independent in the complement of the subpopulation. Against the null hypothesis of a Frank copula, the subpopulation is selected at random and its conditional distribution is forced into a stronger Mallows model. This allows the population margins to remain uniform while possibly restricting the range of the subpopulation.

Figure 10 shows the distribution of p-values of both CSF and CKT tests against 40% Gaussian with  $.12 \leq \rho \leq .13$ . For this range of overall correlation the CKT test holds its level and is conservative for overall correlation  $\rho < .12$ , which is the case here. Nevertheless, it achieves perfect power when the subpopulation  $\rho \geq .125$ , even though its power quickly diminishes to 10% for  $\rho = .12$  in the subpopulation. It also performs better than CSF in this range.



**Figure 10:** P-values against Alternative Mixture of Gaussian copulas with sample size  $n = 1000$  and subpopulation proportion = 40%. Dashed horizontal line at  $\alpha = .05$  indicates the power of a level .05 test.

Under the Mallows alternative,  $n = 1000$  points are generated from a uniform distribution, 400 points are then sampled from a quantile range of  $x$  values and the  $y$  values resorted according to a random draw from a Mallows( $\phi(\tau)$ ) model. Values of  $\tau$  used are .4, .5, and .6. Figure 11 shows the distributions of p-values from the  $\alpha = .05$  level CSF and CKT tests over 100 simulations. The left panel corresponds to the subpopulation being sampled from the full range, while the right panel corresponds to samples between the 20<sup>th</sup> and 80<sup>th</sup> percentiles of  $x$ -values. The CKT test performs much better than the CSF test against these alternatives. The CKT has essentially perfect detection when the subpopulation spans the whole range, and at least 70% power in the 20-80 percentile range. The CSF has no power in either scenario.



**Figure 11:** Distribution of p-values for Mallows alternatives, based on 100 simulations of sample size  $n = 1000$ . Left Panel: Associated subgroup spans the full range of x-values. Right Panel: span is restricted to the 20-to-80<sup>th</sup> percentiles region of x-values in the larger population.

## 5. Concluding Remarks

The ability to detect subpopulations that drive association has the potential of changing the way statistics are used to unveil structures in “Big Data.” Instead of employing extensive model searching with complex interaction, now relatively model-free methods are available to ascertain with precision if there is any simple mixture that better explains monotone association between variables. The CSF and CKT tests achieve this, either working together to screen and confirm or separately to find different forms of the subpopulation that most strongly supports the association.

These tests, however, are formally restricted to different forms of the meaning of “homogeneous association.” Strict legitimacy of the CSF test depends on the assumption of a Gaussian copula underlying the null distribution, whereas the CKT test depends on the assumption of a Frank copula underlying the null distribution. Although there is some evidence of limited robustness, much more work should be done to explore the behavior of these tests under general conditions. For example, both the Gaussian and Frank copulas are radially symmetric; it is unclear how sensitive the tests would be to asymmetric notions of homogeneous association.

The computational efficiency of the CSF test is important because the sample size  $n$  needs to be in the thousands before there is much hope of reliably detecting these subtle but important differences. In contrast with the CSF test, the justification of CKT is a bit more compelling, based on intrinsic association within the subpopulation. We have been using CSF at a liberal  $\alpha = .05$  level as a screening device to reduce the number of pairs of variables to be tested at a more stringent level. Both tests have been employed in the previous section to discover putative gene networks in the liver and the set of chemicals to which each respond.

Detecting heterogeneity of association is a difficult task. Such detection is practical only when the overall association is not too strong, the association in the subpopulation is strong, and the sample size is large. Nevertheless, such scenarios abound. We believe that these new methods will make Statistics ever more relevant in making good sense from Big Data.

### Acknowledgement

This work was partially supported by a grant from the NIH.

### References

- Diaconis, P. and Graham, R.L. (1977). Spearman's footrule as a measure of disarray. *Journal of The Royal Statistical Society Series B - Statistical Methodology*, 39, 262-268.
- Diaconis, P. and Ram, A. (2000). Analysis of systematic scan metropolis algorithms using Iwahori-Hecke algebra techniques. *Michigan Math. J.* **48**, 157-190.
- Frank, M.J. (1979). On the simultaneous associativity of  $F(x, y)$  and  $x + y - F(x, y)$ . *Aequationes Mathematicae* **19**, 194-226.
- Fermanian J.D., Radulovic D. and Wegkamp M. (2004). Weak convergence of empirical copula processes. *Bernoulli* **10**, 847-860.
- Genest, C. (1987). Frank's family of bivariate distributions. *Biometrika*, **74**(3), 549-555.
- Genest, C. and MacKay, J. (1986). The Joy of Copulas: Bivariate Distributions with Uniform Marginals. *American Statistician*,
- Genest, C. and Nešlehová, J. (2014). On tests of radial symmetry for bivariate copulas. *Statistical Papers*, 55.
- Katz, G. (2014). How much do we know about HDL cholesterol? *Clinical Correlations* (<http://www.clinicalcorrelations.org/?p=7298>)
- Lehmann, E.L. and Romano, J.P. (2006). *Testing Statistical Hypotheses*, 3<sup>rd</sup> Edition. Springer: New York.
- Nelsen, R.B. (2006). *An Introduction to Copulas*, 2<sup>nd</sup> Edition. Springer: New York.
- Sampath, S. and Verducci, J. (2013) Detecting the end of agreement between two long ranked lists. *Statistical Analysis and Data Mining*, **6** (6), 458-471.
- Sen, P.K., Salama I.A. and Quade, D. (2011) Spearman's Footrule: Asymptotics in Applications. *Chilean Journal of Statistics*, **2**, 3:20.
- Starr, Shannon (2009). Thermodynamic Limit for the Mallows Model on  $S_n$ . *J. Math. Phys.* **50** 095208
- Voight, B. F. *et al.* (2012) Plasma HDL cholesterol and risk of myocardial infarction: a mendelian randomisation study. *The Lancet*, **380** (9841), 572 - 580.
- Yu, L., Verducci, J. and Blower, P. (2011) The Tau-Path Test for Monotone Association in an Unspecified Subpopulation: Applications to Chemogenomic Data Mining," *Statistical Methodology* **8**, 97-111.