Received: 3 January 2021 | Accepted: 6 August 2021

DOI: 10.1002/qre.3034

SPECIAL ISSUE ARTICLE

WILEY

Bivariate copula monitoring

Andrew Easton[1] | **Okki van Dalen**[2] | **Rainer Goeb**[1] | **Alessandro Di Bucchianico**[2]

[1] Institute of Mathematics, Würzburg University, Würzburg, Germany

[2] Department of Mathematics and Computer Science, Eindhoven University of Technology, Eindhoven, The Netherlands

Correspondence, Andrew Easton, Institute of Mathematics, Würzburg University, Sanderring 2, 97070 Würzburg, Germany.
Email: andrew.easton@mathematik.uni-wuerzburg.de

Funding information
Deutsche Forschungsgemeinschaft

Abstract

The assumption of multivariate normality underlying the Hotelling T^2 chart is often violated for process data. The multivariate dependency structure can be separated from marginals with the help of copula theory, which permits to model association structures beyond the covariance matrix. Copula-based estimation and testing routines have reached maturity regarding a variety of practical applications. We have constructed a rich design matrix for the comparison of the Hotelling T^2 chart with the copula test by Verdier and the copula test by Vuong, which allows for weighting the observations adaptively. Based on the design matrix, we have conducted a large and computationally intensive simulation study. The results show that the copula test by Verdier performs better than Hotelling T^2 in a large variety of out-of-control cases, whereas the weighted Vuong scheme often fails to provide an improvement.

KEYWORDS

copula, multivariate Gaussian distribution, multivariate statistical process control (SPC), phase I, phase II

1 | INTRODUCTION

Copulae have been used on a large scale for dependence modelling in finance,[1] insurance, economics, earth sciences, meteorology, climate research,[2] engineering and chemical process systems engineering.[3] Copula theory enables the modelling and empirical analysis of any practically relevant dependence structure in particular beyond multivariate normal distribution. Further survey of theory and applications is provided by Czado.[4] However, applications of copulae in process monitoring are still marginal as observed by Ren et al.[3] The ample treatment by Czado[4] leaves process monitoring unconsidered.

The application of copulae to process monitoring is pursued in a moderate amount of papers only. Ren et al.[3] use copula-based dependence monitoring for the formulation of methods for the detection of different modes of a given chemical process. Busababodhin and Amphanthong[5] review some existing studies on the use of copulae in process monitoring. The referenced studies concentrate on two subjects: (1) the use of the copula to model serial dependence in data, see Hryniewicz and Szediw[6] as well as Hryniewicz[7] and (2) bivariate copulae to model the dependence of exponentially distributed marginals representing the time between rare events, see Xie et al.[8] and Kuvattna et al.[9] The latter approaches consider shifts only in the exponentially distributed marginals and *not* in the copula describing the dependence structure.

In his monitoring approach, Verdier[10] combines the density level set method for the definition of control limits with copula modelling of a multivariate process. The classical Hotelling T^2 decision rule can then be viewed as a special case

TABLE 1 The confidence limits for the means and correlation of RWE and HeidelbergCement. The scalar Student-*t* confidence levels are 95%

Timeframe	$\hat{\mu}$ lower	$\hat{\mu}$ upper	$\hat{\rho}$ lower	$\hat{\rho}$ upper
2017-12-22 through 2018-05-11	$\begin{pmatrix}0.999\\0.997\end{pmatrix}$	$\begin{pmatrix}1.01\\1.00\end{pmatrix}$	0.084	0.448
2018-05-11 through 2018-09-28	$\begin{pmatrix}0.999\\0.996\end{pmatrix}$	$\begin{pmatrix}1.00\\1.00\end{pmatrix}$	0.059	0.428
2018-09-28 through 2019-02-15	$\begin{pmatrix}0.997\\0.995\end{pmatrix}$	$\begin{pmatrix}1.00\\1.00\end{pmatrix}$	0.203	0.540

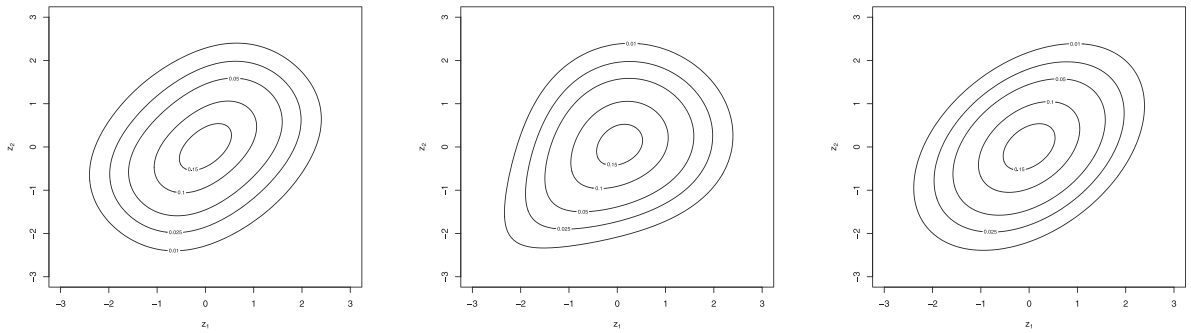

FIGURE 1 Contour plots of copulae fitted to the daily returns of unadjusted prices of RWE and HeidelbergCement

of the density level method under multivariate normal distribution. With respect to out-of-control models, Verdier also restricts attention to shifts in the mean only.

In contrast to all previous authors, we consider process monitoring under a shift in the dependence structure represented by a shift from one copula to another. More specifically, this shift conserves both the marginals and the mean vector. This behaviour occurs particularly in financial return time series. One example among many is the dependence structure of the stocks of RWE and HeidelbergCement during three consecutive blocks of 100 trading days from the end of 2017 through early 2019.

When analysing the dependence structure of the daily returns of the daily unadjusted prices[11] of RWE and HeidelbergCement over 300 trading days from 22 December 2017 through 15 February 2019 in three blocks of 100 consecutive trading days the confidence intervals for the mean vector and the correlation all overlap as shown in Table 1. So from this point of view the contour plots of the estimated dependence copulae in Figure 1 should intuitively be structurally similar. However, the three corresponding copula estimates shown in Figure 1 indicate that there may nevertheless be shifts in tail dependence when going from one block of 100 days to the next. Note the stronger lower tail dependence on the lower left quadrant in the middle plot, which does not appear in the plots on the left or the right. Furthermore, there appears to be no upper tail dependence between the two stocks in any of the three estimated copulae, as evidenced by the upper right quadrants.

When considering all pairs of DAX-30 stocks from 28 November 2014 through 28 November 2019 then there are in total 5234 blocks of 100 trading days. During the analysis of a given pair of stocks, bivariate copula models are estimated for each block of 100 days. Then for *consecutive* blocks, the models are compared. Overall, in 302 cases, the correlation and mean of two consecutive models were the same while the copula families and with them the estimated dependence structure differed. Therefore, over 5% of the observed dependence changes between two consecutive blocks preserve the mean and the correlation.

When analysing dependence between two or more variables, it is necessary to model such changes in the dependence structure. However, in cases such as this one considering mean and correlation only is not sufficient.

This paper describes and analyses a flexible approach to monitoring bivariate dependence and in particular tail-dependence. The study is organised as follows: Section 2 explains the general rationale of copula theory in process monitoring. Section 4 outlines two copula-based monitoring charts. Section 5 lays out the design matrix underlying both simulation studies. Section 6 presents the results of the two simulation studies put forth in this paper.

TABLE 2 The design matrix. Please note that $\widetilde{c}(\theta) = \boldsymbol{u} \mapsto c(\boldsymbol{u}; \theta)$, where $\widetilde{c} \in \{Gaussian, Gumbel, Clayton, Joe, Frank\}$ and c is the corresponding copula density in the notation used in Section 2

Phase II ρ	Phase I	Phase II	Vuong Phase II weight parameter
0.8572962	Gaussian(**0.8572962**)	Gumbel(**3**)	$\lambda \in \{\text{NA}, 0.1, 0.2, 0.3, 0.4, 0.5\}$
0.9714609	(See above)	Gumbel(7)	(See above)
0.0010472	(See above)	Gumbel(1)	(see above)
0.7739141	Gaussian(**0.7739141**)	Clayton(**3**)	(See above)
0.9032978	(See above)	Clayton(7)	(See above)
0.4990976	(See above)	Clayton(1)	(See above)
0.6993231	Gaussian(**0.6993231**)	Joe(**3**)	(See above)
0.8886440	(See above)	Joe(7)	(See above)
0.0005751	(See above)	Joe(1)	(See above)
0.4251543	Gaussian(**0.4251543**)	Frank(**3**)	(See above)
0.7235313	(see above)	Frank(7)	(See above)
0.1571861	(See above)	Frank(1)	(See above)
0.6	Gaussian(**0.3**)	Gaussian(0.6)	(See above)
0.9	Gaussian(**0.6**)	Gaussian(0.9)	(See above)
0.6	Gaussian(**0.9**)	Gaussian(0.6)	(See above)

The bold values signify the most advantageous ARLs. In this case the smallest ones.

2 | RATIONALE OF COPULA MONITORING

The Hotelling T^2 chart is widely used for multivariate process monitoring. Its control limits depend on the Phase I mean vector $\boldsymbol{\mu}$ and on the Phase I covariance matrix $\boldsymbol{\Sigma}$. These two parameters miss some of the dependence structure a multivariate distribution may offer as outlined above for the daily returns of RWE and HeidelbergCement. How does the Hotelling T^2 chart fare in situations like these?

In comparison to Figure 1, this paper considers more articulate examples of changes in the dependence structure. It assumes that all marginal distributions are standard normal and only considers copulae that depend on one parameter, which shall be called θ. Specifically, the Gaussian, Gumbel, Clayton, Joe and Frank copula families were chosen. Note that the correlation ρ is strictly increasing in θ and in the case of the Gaussian copula $\rho = \theta$. Additionally, the changes in tail dependence are chosen in a way that makes them somewhat more pronounced than in Figure 1, in order to facilitate detection. Also a few more variations or changes in dependence structure are considered in order to get a good overview of different underlying changes, which may occur during process monitoring. Nevertheless, these simplifications and slight adjustments yield worthwhile results by providing valuable guidelines.

The targeted examples are the rows in Figures 2 and 3. They show contour plots of bivariate distributions with standard normal marginals meaning that $\boldsymbol{\mu} = \mathbf{0}$ and $\sigma_1^2 = \sigma_2^2 = 1$. The correlations ρ range from about 0.43 to approximately 0.86 and are specified for each row individually. As the variances and the correlation are the same for each row, the corresponding covariance matrices are also the same. The plots on the left hand side show bivariate normal distributions where the dependence only depends on ρ. However, the plots on the right hand side are *not* constructed from the Gaussian copula, meaning they describe non-normal bivariate distributions. Therefore, these examples show that some aspects of dependence structure are described by the correlation ρ while others are not.

As to the parameters, the copula parameter $\theta = 3$ was chosen as a point of reference. Lower dependence was modelled with $\theta = 1$ and higher dependence with $\theta = 7$, with both values exhibiting visible differences to the reference point. As the plots clearly show there may be contexts in which it is undesirable to consider two densities on the same row to be equivalent, ensuring that these examples capture the thrust of the example of RWE and HeidelbergCement. These cases are later also featured in our design matrix in Table 2.

Naturally, this implies that only estimating the mean when applying the Hotelling T^2 chart cannot differentiate the pairs of cases given by the rows in Figures 2 and 3, but instead consider them as equivalent, because $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are estimated to be the same. Therefore, two additional approaches for monitoring *non*-normal multivariate dependence structure beyond the covariance matrix will be presented, which are both based on copula theory. The results help clarify the applicability of these two approaches as well as their relationship to the widely used Hotelling T^2 chart.

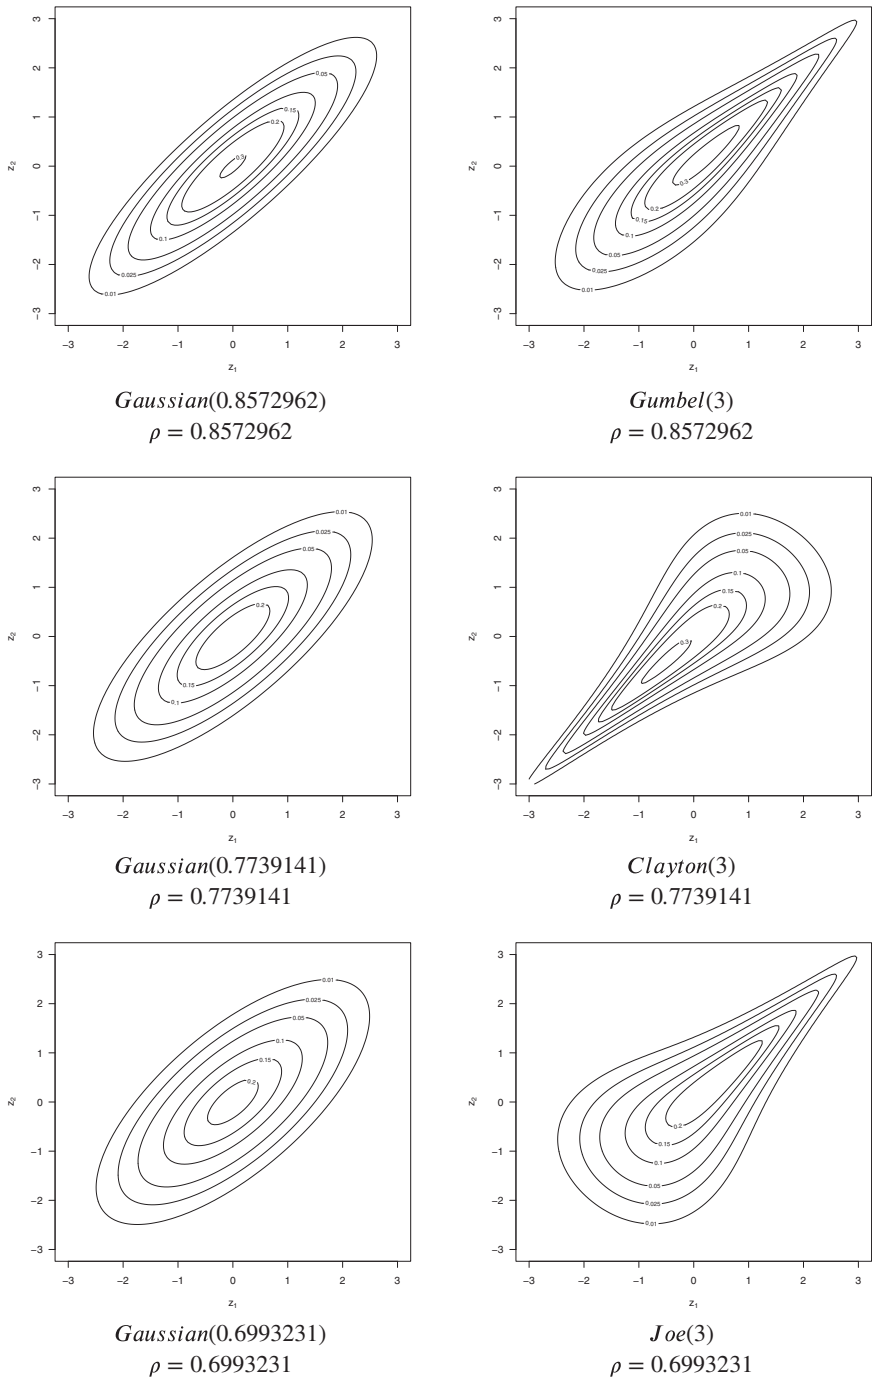

FIGURE 2 Contour plots of bivariate densities with univariate standard normal marginals. The dependence structure is given by the respective copulae

3 | COPULA THEORY FOR PROCESS MONITORING

Copula theory is motivated by decomposing a multivariate distribution into its marginals and a factor describing the dependency structure based on Sklar's Theorem,[12] see the exposition by Czado.[4] While Nelsen[13] offers a general proof of the theorem for any type of distribution, the focus here lies on distributions of absolutely continuous type where both the CDF and PDF can be decomposed *uniquely*. Let the random vector $\boldsymbol{X}$ follows an absolutely continuous distribution. Then its PDF can be uniquely decomposed in the following way:

$$f_{\boldsymbol{X}}(\boldsymbol{x};\boldsymbol{\theta}) = c\big(F_{X_1}(x_1;\boldsymbol{\theta}),\dots,F_{X_N}(x_N;\boldsymbol{\theta});\boldsymbol{\theta}\big)\cdot f_{X_1}(x_1;\boldsymbol{\theta})\cdots f_{X_N}(x_N;\boldsymbol{\theta})$$

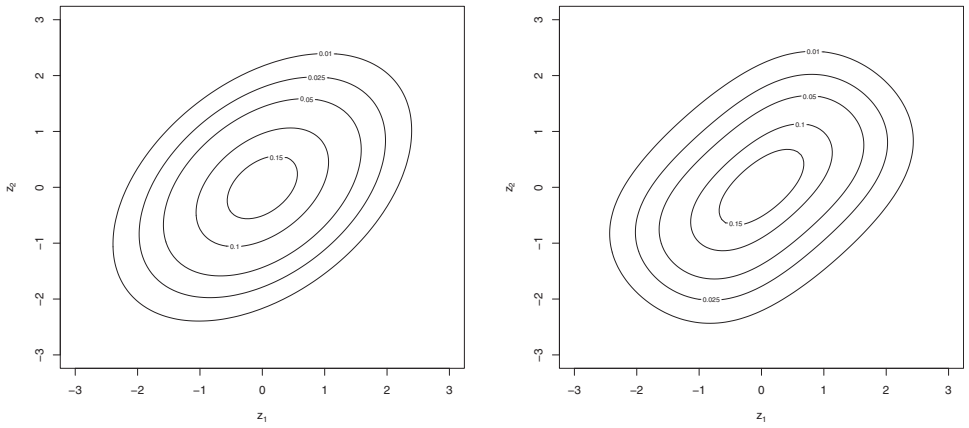

$Gaussian(0.4251543)$ $\rho = 0.4251543$ $Frank(3)$ $\rho = 0.4251543$

FIGURE 3 Contour plots of bivariate densities with univariate standard normal marginals. The dependence structure is given by the respective copulae

where the function $c : [0;1]^N \times \Theta \longrightarrow \mathbb{R}_0^+$ is called the copula density, N is the dimension of the random vector $\boldsymbol{X}$ and Θ denotes the parameter space. The equality means that the copula density c describes all of the dependence structure inherent in the arbitrary absolutely continuous density on the left hand side, as the N remaining factors $f_{X_i}(x_i;\boldsymbol{\theta}), i = 1,\dots,N$ on the right hand side are all univariate marginal densities. Furthermore, as this decomposition is guaranteed to exist, this equation also permits to construct an arbitrary absolutely continuous distribution's density by choosing a copula density and univariate marginals.

As a first example a non-degenerate bivariate normal density is decomposed

$$\varphi_X(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma}) = c(\Phi_1(x_1),\Phi_2(x_2);\rho)\cdot\varphi_1\left(x_1;\mu_1,\sigma_1^2\right)\cdot\varphi_2\left(x_2|\mu_2,\sigma_2^2\right)$$

where

$$\boldsymbol{x} = (x_1,x_2)^\top, \boldsymbol{\mu} = (\mu_1,\mu_2)^\top, \boldsymbol{\Sigma} = \begin{pmatrix}\sigma_1^2 & \sigma_{12}\\ \sigma_{12} & \sigma_2^2\end{pmatrix}$$

and $\rho = \sigma_{12} \,/\, (\sigma_1\sigma_2)$.

Note, that the unique copula resulting from the decomposition of the multivariate normal distribution according to Sklar's theorem is called the Gaussian copula with copula density

$$Gaussian(u_1,u_2|\rho) = \frac{1}{\varphi(y_1)\varphi(y_2)}\frac{1}{\sqrt{1-\rho^2}}\exp\left\{-\frac{\rho^2\left(y_1^2+y_2^2\right)-2\rho y_1 y_2}{2(1-\rho^2)}\right\}$$

where $y_1 = \Phi^{-1}(u_1)$, $y_2 = \Phi^{-1}(u_2)$ and Φ, φ are the distribution function and density of the univariate *standard* normal distribution, respectively. This immediately leads to two crucial observations: First, the copula density only depends on the correlation ρ but neither on the means μ_1,μ_2 nor the variances σ_1^2,σ_2^2. Recall, that multivariate normality implies that two marginals are independent if and only if they are uncorrelated. Second, due to the uniqueness of the density's decomposition, an absolutely continuous distribution is bivariate normal if and only if its copula density is the density of the Gaussian copula and the marginals are univariate normal distributions.

The latter leads to a very useful intuition concerning the nature of the multivariate normal distribution in relation to other multivariate distributions. It is possible to *violate* the assumption of multivariate normality by choosing at least one non-Gaussian *marginal* or by choosing a *copula* that is not the Gaussian copula. This permits to purposefully violate the normality assumption of the Hotelling T^2 chart.

For instance, changing one marginal to follow an exponential distribution as outlined in Example 1.15 of Czado[4] gives a good intuition as to what it means for a non-normal multivariate distribution to have the same copula and hence the *same dependence structure* as a bivariate normal distribution even though one marginal is non-normal. This sameness of dependence structure is defined in a strictly formal way. Moreover, the Gaussian copula is an example of a bivariate dependence structure that is fully described by the correlation.

In this paper, the focus lies on preserving the marginals and changing the copula family, while preserving the correlation. Note, that in general changing the copula family may affect the entire dependence structure. By changing the copula family and choosing a parameter, which preserves the correlation in spite of the change, examples are constructed where the dependence is affected in a way that cannot be described by only considering the correlation. Even though this procedure relies on choosing the parameters carefully, the example of RWE and HeidelbergCement demonstrates the practical interest in these kinds of phenomena. For completeness, we will also consider examples that change both the copula family and the correlation. Note, that the marginals remain constant and therefore unaffected by the changes. Section 5 will demonstrate how these crucial aspects of Figures 2 and 3 are reflected in our design matrix while adding both the above mentioned structural changes that preserve correlation and the ones that changes correlation. The latter provide insight into simultaneous changes of the dependence copula and the correlation.

The copulae in Figure 4 are additional copula families available for estimation in the second simulation study. Summarily, the utilised copula densities are given in the following and example contours are plotted in Figures 2–4:

$$Gumbel(u_1,u_2|\theta) = \frac{\partial^2}{\partial u_1 \partial u_2} exp\left[-\{(-\ln u_1)^\theta + (-\ln u_2)^\theta\}^{\frac{1}{\theta}}\right]$$

$$Clayton(u_1,u_2|\theta) = (1+\theta)(u_1^{-\theta}+u_2^{-\theta}-1)^{\frac{1+2\theta}{\theta}}(u_1u_2)^{-(\theta+1)}$$

$$Joe(u_1,u_2|\theta) = \frac{\partial^2}{\partial u_1 \partial u_2} 1-\left((1-u_1)^\theta + (1-u_2)^\theta - (1-u_1)^\theta(1-u_2)^\theta\right)^{\frac{1}{\theta}}$$

$$Frank(u_1,u_2|\theta) = \frac{\partial^2}{\partial u_1 \partial u_2} - \frac{1}{\theta}\ln\left(\frac{1}{1-e^{-\theta}}[(1-e^{-\theta}) - (1-e^{-\theta u_1})(1-e^{-\theta u_2})]\right)$$

$$Student-t(u_1,u_2|\nu,\rho) = \frac{t(T_\nu^{-1}(v_1),T_\nu^{-1}(v_2)|\nu,\rho)}{t_\nu(T_\nu^{-1}(v_1))t_\nu(T_\nu^{-1}(v_2))}$$

$$BB1(u_1,u_2|\theta,\delta) = \frac{\partial^2}{\partial u_1 \partial u_2}\eta(\eta^{-1}(u)+\eta^{-1}(v)), \text{ where } \eta(s) = \eta_{\theta,\delta}(s) = (1+s^{\frac{1}{\delta}})^{-\frac{1}{\theta}}$$

$$BB6(u_1,u_2|\theta,\delta) = \frac{\partial^2}{\partial u_1 \partial u_2} 1-\left(1-\exp\left(-((-\ln(1-(1-u_1)^\theta))^\delta + (-\ln(1-(1-u_2)^\theta))^\delta)^{\frac{1}{\delta}}\right)\right)^{\frac{1}{\theta}}$$

$$BB7(u_1,u_2|\theta,\delta) = \frac{\partial^2}{\partial u_1 \partial u_2}\eta(\eta^{-1}(u)+\eta^{-1}(v)), \text{ where } \eta(s) = \eta_{\theta,\delta}(s) = 1-\left[1-(1+s^{-\frac{1}{\delta}})\right]^{\frac{1}{\theta}}$$

$$BB8(u_1,u_2|\theta,\delta) = \frac{\partial^2}{\partial u_1 \partial u_2}\frac{1}{\delta}\cdot\left(1-\left(\frac{1}{1-(1-\delta)^\theta}(1-(1-\delta u_1)^\theta)(1-(1-\delta u_2)^\theta)\right)^{\frac{1}{\theta}}\right)$$

In summary, copula theory provides a two-pronged approach. First, it permits to examine and evaluate established monitoring charts in situations where a selection of their assumptions are violated. For example, the Hotelling T^2 chart can be analysed for cases where the underlying distribution is *non*-normal in a precisely controlled way. This permits to examine how different violations of the underlying assumptions affect a monitoring method's performance in practice. Second, as demonstrated in Verdier,[10] copula theory also gives rise to new methodologies. The types of structural

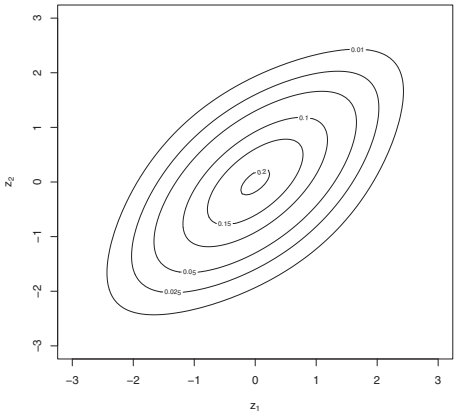

$Student - t(15, 0.6)$
$\rho = 0.6$

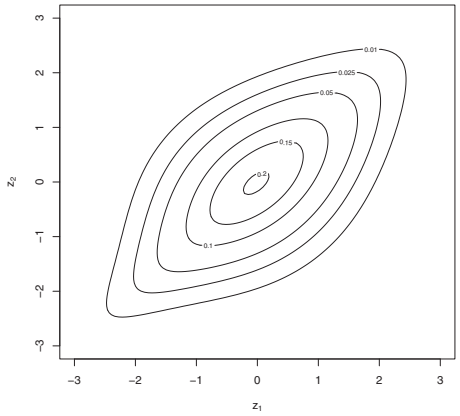

$BB1(0.35, 1.39)$
$\rho = 0.55$

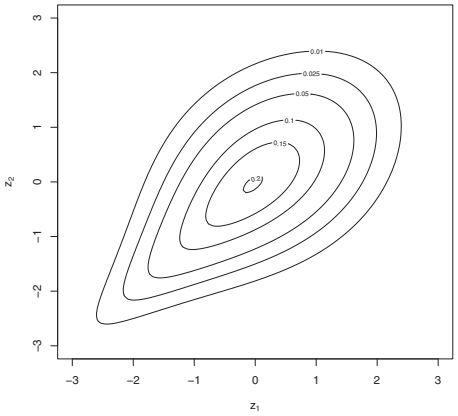

$BB6(1, 1.6)$
$\rho = 0.53$

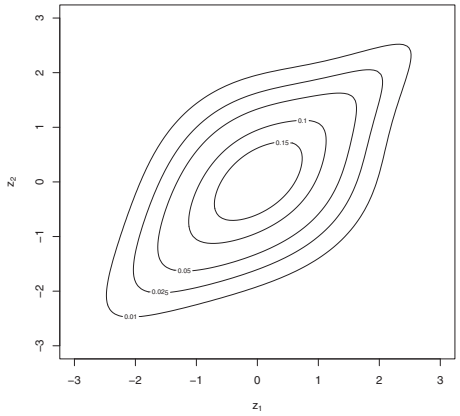

$BB7(1.47, 0.65)$
$\rho = 0.52$

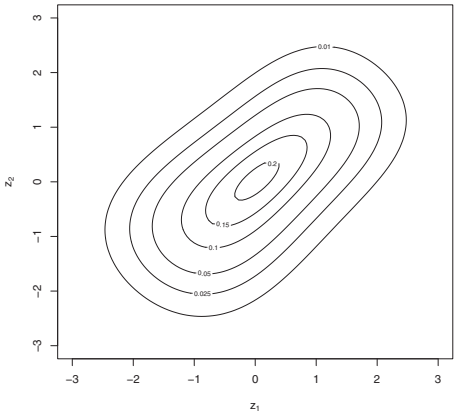

$BB8(8, 0.43)$
$\rho = 0.58$

FIGURE 4 Contour plots of bivariate densities with univariate standard normal marginals. The dependence structure is given by the respective copulae

dependence illustrated in Figures 1–3 indicate that copula theory can provide monitoring charts, which allow to consider dependence structures other than the covariance matrix.

4 | COPULA MONITORING

We analyse two approaches to process monitoring based on copula theory. The first is the approach by Verdier[10]. The second is constructed by combining the Vuong[14] test statistic with weighted copula estimation.

4.1 | Hotelling's T^2 chart

The Hotelling T^2 chart is used with fixed $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, which are assumed to be known from phase I, see Mason and Young.[15]

4.2 | Verdier's monitoring chart

The monitoring chart proposed by Verdier[10] uses density level set estimation to obtain a multivariate tolerance region based on Baíllo and Cuevas.[16] The following outline follows Verdier's paper closely. This general approach is implemented by estimating copula models for Phase I allowing Monte Carlo estimations of the tolerance regions and computation of the density values of interest for large Monte Carlo sample sizes. The underlying and in practice unknown Phase I density is denoted as f. The tolerance region D depends on the false-alarm rate α. The desirable properties are the following:

$$D(c_\alpha) = \left\{x \in \mathbb{R}^d | f(x) \geq c_\alpha\right\} \tag{1}$$

where c_α satisfies

$$P(X \in D(c_\alpha)) \ = \int_{D(c_\alpha)} f(x)\, dx = 1 - \alpha. \tag{2}$$

In an empirical analysis, f is replaced with the estimated $\widehat{f}$. Verdier[10] demonstrates that the properties (1) and (2) are compatible with his proposed estimators for both f and the level sets of f.

The illustration in Figure 5 is univariate, even though the Verdier chart is applied to the multivariate Phase I density in the simulation study. The goal is to find one or two level sets that describe very unlikely or very likely Phase I realisations. In practice, a Monte Carlo method is used to estimate these level sets. When transferred back to the density's domain, the level sets marked in red correspond the measurements, which trigger a signal.

Verdier's asymptotic analysis hinges on three key assumptions

1. The density f is of class C^2 with a bounded Hessian matrix, and $f(x) \longrightarrow 0$ as $\|x\| \longrightarrow \infty$.
2. $\mathcal{H}_0$ has Lebesgue content 0
3. $\lambda(\{f = c\}) = 0$ for every $c > 0$.

where

$$\mathcal{H}_0 = \left\{ c \in \left(0; \sup_{x \in \mathbb{R}^d} f(x)\right) | \inf_{\{x | f(x) = c\}} \|\nabla f(x)\| = 0 \right\}$$

following Cadre et al.[17]

Verdier's approach permits to construct an upper control limit, a lower control limit or both limits at the same time. Mühlig's[18] analysis demonstrates that choosing both an upper and lower control limit leads to more sensitive change detection compared to only constructing a lower control limit. The double-sided tolerance region is constructed as

$$D(\ell_\alpha, u_\alpha) = \left\{x \in \mathbb{R}^d | \, \ell_\alpha \leq f(x) \leq u_\alpha\right\} \tag{3}$$

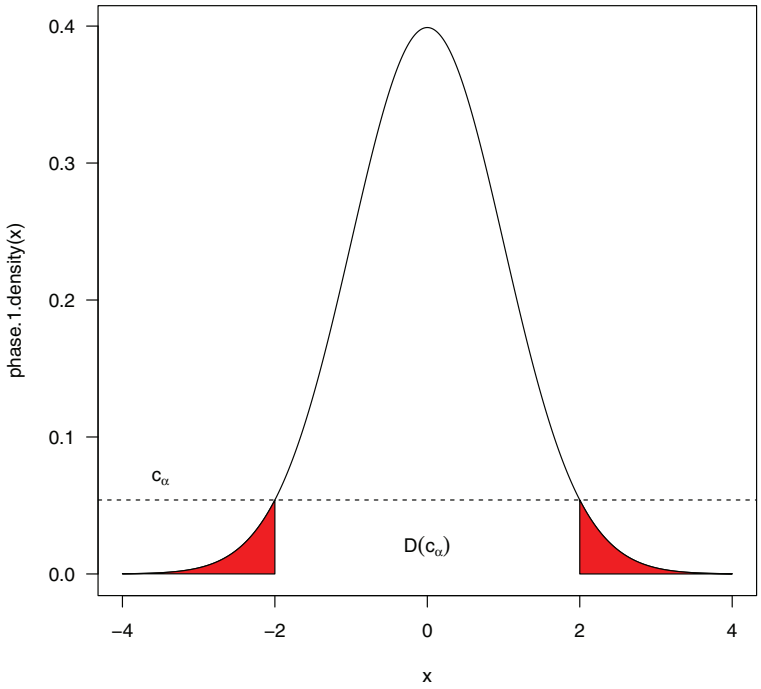

FIGURE 5 Univariate illustration of the lower control limit constructed via level sets by Verdier's[10] approach. The lower control-limit is set to c_α leading to tolerance region $D(c_\alpha) = [-2; 2]$

where again

$$P(X \in D(\ell_\alpha, u_\alpha)) = \int_{D(\ell_\alpha, u_\alpha)} f(x)\, dx = 1 - \alpha. \tag{4}$$

In the simulation discussed in Section 4.2 study ℓ_α and u_α in Equations (3) and (4) are chosen such that for

$$U(u_\alpha) = \left\{x \in \mathbb{R}^d \,|\, u_\alpha < f(x)\right\},$$
$$L(\ell_\alpha) = \left\{x \in \mathbb{R}^d \,|\, f(x) < \ell_\alpha\right\}$$

it holds that

$$P(X \in U(u_\alpha)) = \int_{U(u_\alpha)} f(x)\, dx = \alpha/2,$$
$$P(X \in L(\ell_\alpha)) = \int_{L(\ell_\alpha)} f(x)\, dx = \alpha/2.$$

The choice of picking $\alpha/2$ for both the lower and upper bound is made in accordance with Mühlig.[18]

4.3 | Weighted copula estimation and the Vuong test

Our second monitoring approach utilises Vuong's[14] test statistic ν applied to weighted copula estimates. The Phase I PDF estimation uses no weights. All subsequent samples containing at least one phase II observation are weighted. The weights emphasise the newest observations over the older ones in order to make the estimation procedure adaptive to changes in the copula structure. The use of weights is suggested by Nagler.[19] Though not used in Vuong's[14] original approach, weights can be integrated perfectly into his asymptotic analysis of the test statistic.

Vuong's[14] two-sided likelihood ratio test is based on the null hypothesis that two non-nested copulae are to be considered the same on a given sample. The null hypothesis can be rejected in favour of one of the two tested copulae, thereby choosing one copula model over the other.

The theoretical background of Vuong's[14] test statistic depends on two assumptions: (1) asymptotic of the test statistic holding for large sample size N and (2) non-nestedness of models. With respect to (1), Vuong[14] does not provide any constructive inequalities for the accuracy of the approximation as a function of N. In analogy to the experience with other asymptotics, the application on small sample sizes is not unreasonable on principle. With respect to (2), Brechmann et al.[20] demonstrate for large sample sizes, that the violation of non-nestedness is not a serious obstacle to a valid application of Vuong's test. Hence the application of the Vuong test to overlapping samples is of interest. The use of weighted likelihood as a basis of Vuong's test is suggested by Schepsmeier et al.[21] In general, weighted likelihood is considered in the literature in tests related to Vuong's test, confer Amisano and Giacomini.[22] In the following, we call this setting the weighted Vuong test.

It is prudent to consider a large variety of parametric copula families in order to allow the Phase II copula estimation to be as sensitive as possible. Since copula families were included where the inversion of Kendall's tau is impossible, ML estimation was used throughout. The following copula families were used *Gaussian*, *Gumbel*, *Clayton*, *Joe*, *Frank*, *Student-t*, *BB1*, *BB6*, *BB7*, *BB8*.

Weights are infiltrated into the MLE of copula density $\widehat{c}$ and the parameter vector $\widehat{\theta}$ according to the following equation:

$$(\boldsymbol{u}, \boldsymbol{w}) \longmapsto (\widehat{c}, \widehat{\theta}) := \arg\max_{c,\theta} \sum_{i=1}^{N} w_i \cdot \ln c(u_i|\theta), \tag{5}$$

where $\boldsymbol{u} \in \mathbb{R}^N$ is the sample vector, $\boldsymbol{w} \in \mathbb{R}^N$ is the weight vector, $N \in \mathbb{N}$ is the sample size, and the parameters must lie in the parameter space $\theta \in \boldsymbol{\Theta}$. Using Equation (5) on two samples $\boldsymbol{u}_1, \boldsymbol{u}_2$ corresponding weight vectors $\boldsymbol{w}_1, \boldsymbol{w}_2$ yields $\widehat{c_1}, \widehat{\theta_1}$ and $\widehat{c_2}, \widehat{\theta_2}$. Vuong's test-statistc ν has three input quantities: two alternative copula densities $\widehat{c_1}, \widehat{\theta_1}$ and $\widehat{c_2}, \widehat{\theta_2}$ as well as a data vector, here the data vector $\boldsymbol{u}_2$ from the second model.

$$\begin{aligned} H_0 &: \mathrm{E}[m_1] = \mathrm{E}[m_2] = 0 \\ m_i &:= \log\left[\frac{\widehat{c_1}(u_{2,i}|\widehat{\theta_1})}{\widehat{c_2}(u_{2,i}|\widehat{\theta_2})}\right] \\ \nu &:= \frac{\frac{1}{\sqrt{N}}\sum_{i=1}^{N} m_i}{\sqrt{\sum_{i=1}^{N}(m_i - \overline{m})^2}} \end{aligned}$$

The test is two-sided—denoting the significance level als α:

$$\begin{aligned} \nu > \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) &\text{ leads to rejecting } H_0 \text{ in favour of model } \widehat{c_1}, \widehat{\theta_1} \\ \nu < \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) &\text{ leads to rejecting } H_0 \text{ in favour of model } \widehat{c_2}, \widehat{\theta_2} \\ &\text{while } |\nu| \leq \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) \text{ means } \textit{not} \text{ rejecting } H_0. \end{aligned}$$

The use of weights makes the estimation process adapted, which permits to analyse mixed samples containing both Phase I data and Phase II data, while the importance of the Phase II data can be emphasised. This is especially useful when there are *too few Phase II samples* available for estimating a copula on a pure Phase-II-sample.

The samples comprised purely of Phase I data are *always* unweighted $\boldsymbol{w}_1 = (1, \ldots, 1)^\top$, because all of the data are interpreted to have been previously screened and therefore be of relevance irrespective of the samples' time indices in stark contrast to Phase II data. The latter has the property that newer samples are interpreted to be of higher importance than older samples, because being out-of-control should be detected as soon as possible.

4.4 | An illustration of the weighted Vuong methodology

Figure 6 illustrates the weighted Vuong methodology. The illustration shows seven consecutive data points across time. Samples of size $N = 4$ are taken consecutively at each time point. Subsequent samples have an overlap of three time points.

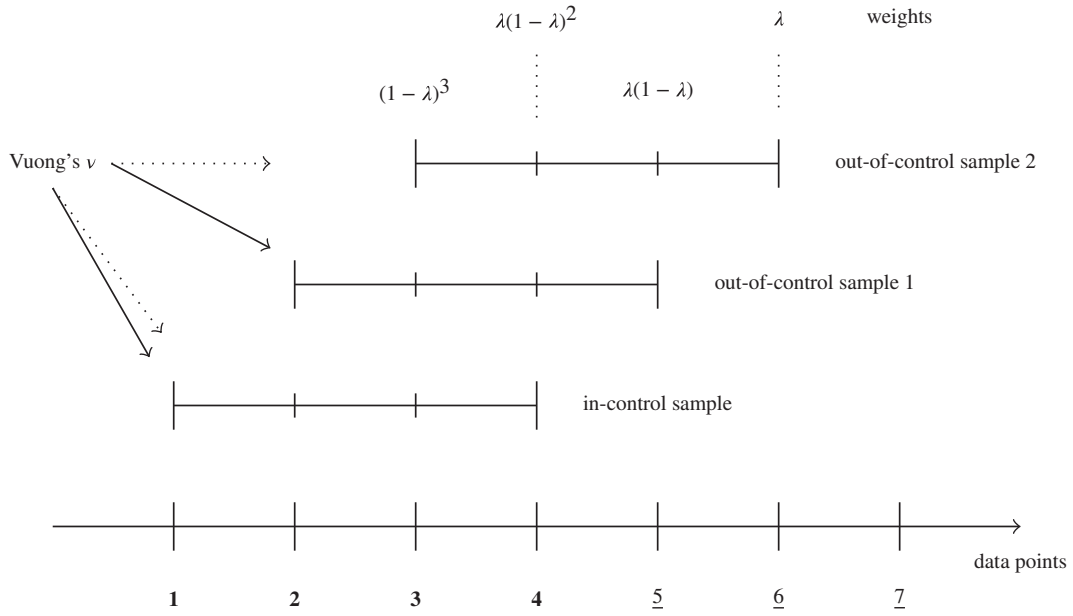

FIGURE 6 Illustration: Vuong's[14] ν and iterated weighted Phase II estimation, see Section 4.4

Let the in-control case be made up of the first four data points. In the figure, all Phase II data points are out-of-control. The first out-of-control sample at time 5 is comprised of the youngest three in-control data points and the first out-of-control data point. The second out-of-control sample at time 6 is comprised of the youngest two in-control data points and the oldest two out-of-control data points. The top of the figure shows the four weights of the four observations in the out-of-control samples. Pure Phase I samples are always unweighted.

One copula model is estimated for every batch of out-of-control samples, while the in-control model is assumed to be known. For the out-of-control models, the marginals are assumed to be known, but not the copula. Note, that this is a very strong assumption. To make the monitoring process adaptive, it seems reasonable to also consider weighted copula estimation.

The weights for the out-of-control batches are inspired by exponential smoothing. The weighting parameter is called λ and the sequence of weights is generated as in exponential smoothing. The unweighted Phase II estimation case is analysed for comparison and is denoted by $\lambda = NA$. The in-control batch is never weighted, because all in-control samples are interpreted to be of equal importance. Note, that the first out-of-control sample is also weighted, even though the illustration only shows weights for the second.

The Phase I model is iteratively tested against Phase II models until the Vuong test refutes its null hypothesis in favour of the Phase II model. This is subsequently interpreted as a signal that the process is out-of-control.

5 | DESIGN MATRIX OF SIMULATION STUDIES

As Sections 1 and 2 discussed that there is a reason to specifically examine changes in the copula density of a given absolutely continuous bivariate distribution. In the following, 12 cases are laid out in which the bivariate copula density is changed while the given standard normal marginals are preserved. Furthermore, the copula density for copula parameter $\theta = 3$ is changed in such a way as to preserve both the joint mean $\boldsymbol{\mu}$, the marginal variances and the correlation ρ, meaning that the change in dependence structure must lie beyond the covariance matrix. The univariate marginals are always standard normal.

In our simulation studies, copulae from the Gumbel, Clayton, Joe and Frank copula families are substituted for the Gaussian copula. All four of these are single parameter copula families. To quantify their effects on the joint mean $\boldsymbol{\mu}$, four samples of size 1024^2 are generated and analysed. The results show that the Student-t confidence intervals for the joint means contain $\mathbf{0}$ for the copula parameters $\theta \in \{1, 3, 7\}$ for each of the four copula families under consideration. This is interpreted as these copulae preserving the joint mean. Furthermore, the same samples are used to calculate the

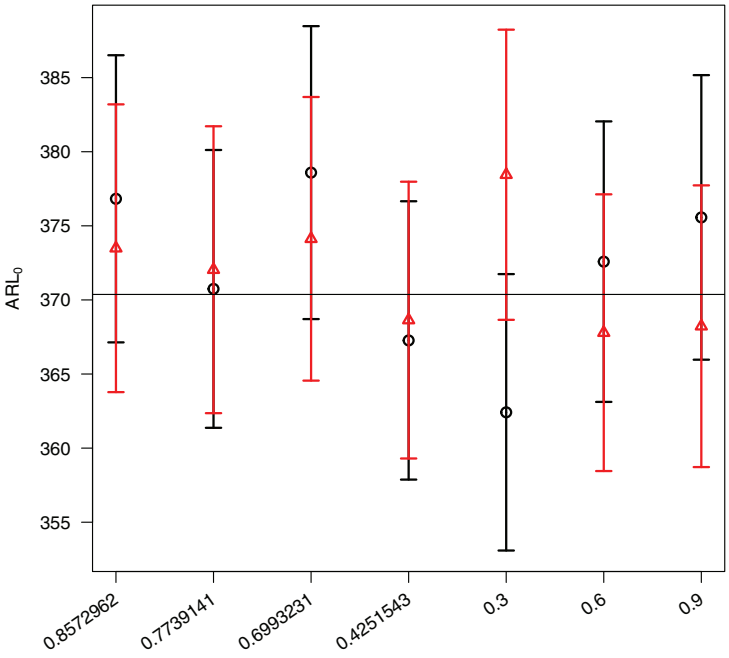

FIGURE 7 In-control average run-lengths ARL_0 for T^2 in black circles and **Verdier** in red triangles. Confidence level 99%. Confer Table 2. The black horizontal line is the target ARL_0 for $\alpha = 0.0027$

resulting correlations. All four families' correlations are monotone in θ, so $\theta = 3$ is chosen as a baseline for the four Pearson correlations, which are to be preserved. Each of the resulting four copulae is paired with a bivariate Gaussian copula of the same correlation, as seen in Table 2, because this permits to preserve $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ when starting with the Gaussian copula and a given correlation and then switching to the associated alternative copula, all else and in the particular the marginals being equal. Note, that this also guarantees Phase I to follow a bivariate normal distribution.

In order to analyse simultaneous changes in the copula and the correlation, the copula parameter θ of the associated copula family is set to 7 for a higher correlation and to 1 for a lower correlation. This yields 12 entries in the design matrix in which Phase I is bivariate normal whereas Phase II is not while the standard normal marginals and the mean vector $\boldsymbol{\mu}$ are preserved.

Finally, three bivariate-normal-to-bivariate-normal cases are added for comparison. These three cases preserve the standard normal marginals as well as the copula family, see Section 2, and therefore necessarily change the correlation parameter of the Phase I Gaussian copula in order to construct an out-of-control Phase II model.

Both simulation studies presented in this paper analyse the estimated average run-lengths ARL of the in-control and the out-of-control cases by averaging 10,000 sample run-lengths. For the Verdier approach, the sample size was chosen to be $N = 10$. For the Vuong approach, the sample size was chosen to be $N = 10$. The in-control chart calibration for analysing the Verdier chart targets an average in-control run-length of $1/\alpha = 1/0.00270 \approx 370$ while the calibration for the Vuong test with weighted or unweighted Phase II samples is calibrated to an in-control average run-length of 365 ± 2.5, which corresponds to $\alpha \in [0.00272; 0.00276]$. The simulation size of 10,000 sample run-lengths is chosen in order to acquire useful confidence intervals for the Vuong approach.

In summary, the design matrix analyses mean and covariance preserving dependence changes away from bivariate normality in four cases, the changing of correlation in three bivariate normal cases and the change of both the dependence and the correlation in eight cases.

6 | RESULTS OF SIMULATION STUDY

6.1 | Verdier's monitoring chart

We consider the results pertaining to the Verdier chart. All analysed cases are calibrated to an in-control ARL_0 corresponding to $\alpha = 0.0027 \approx 1/370$. The in-control model calibration is validated to guarantee comparability with the help of a simulation. The calibration results including their confidence intervals are displayed in Figure 7. All models are successfully calibrated.

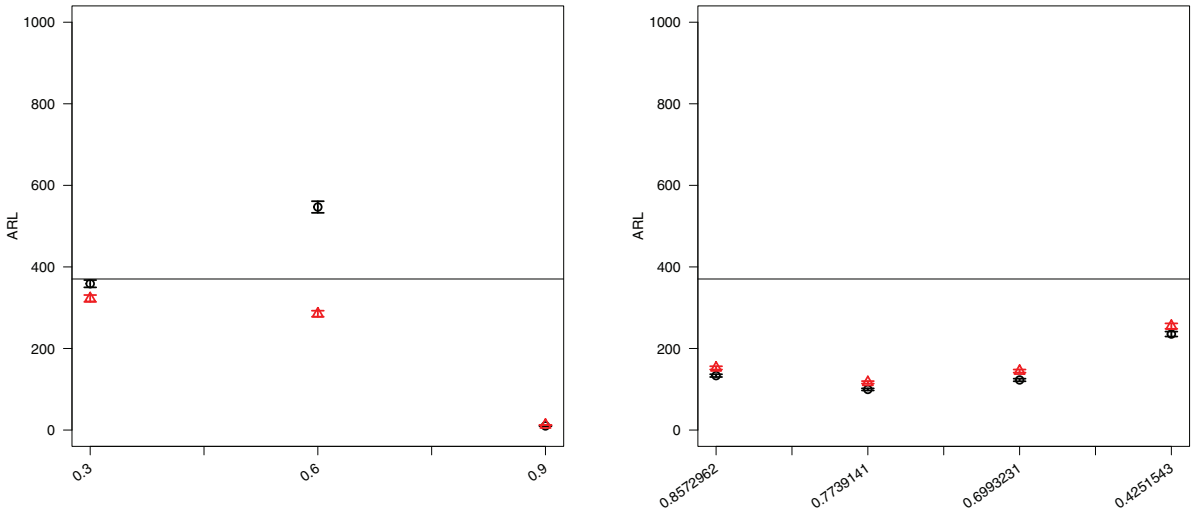

Phase II biv. normal

Phase II non-normal and $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are preserved

FIGURE 8 Out-of-control average run-lengths *ARL* for T^2 in black circles and **Verdier** in red triangles. Student-*t* confidence level 99% computed on 10,000 samples. The black horizontal line is the target ARL_0 for $\alpha = 0.0027$.

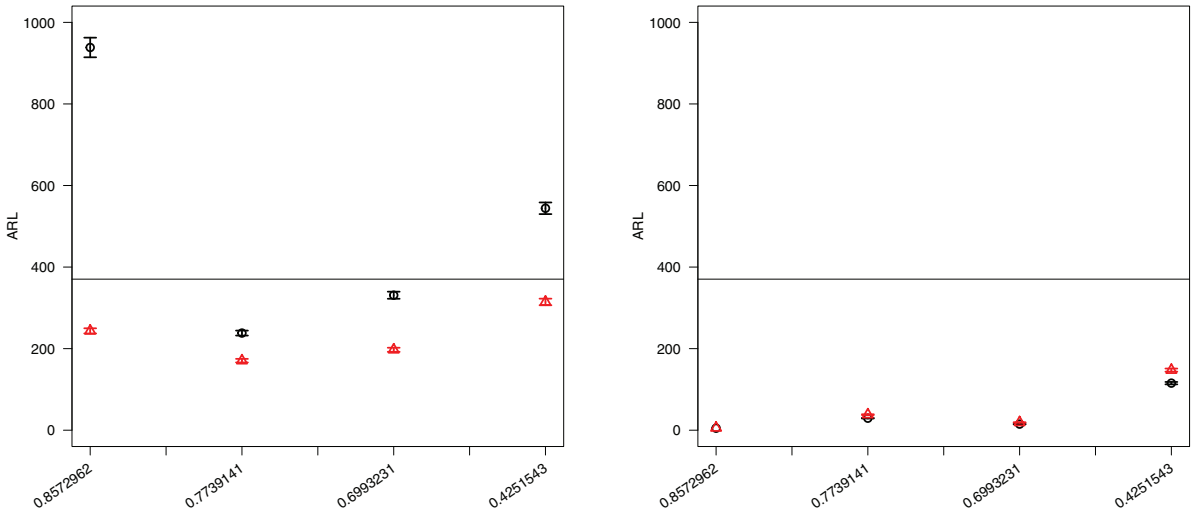

Phase II non-normal, ρ raised and μ preserved.

Phase II non-normal, ρ lowered and μ preserved.

FIGURE 9 Out-of-control average run-lengths *ARL* for T^2 in black circles and **Verdier** in red triangles. Student-*t* confidence level 99% computed on 10,000 samples. The black horizontal line is the target ARL_0 for $\alpha = 0.0027$

The out-of-control results are illustrated by Figures 8 and 9. First, three of the Hotelling T^2 chart's out-of-control *ARL*s excessively overshoot the target ARL_0. All three of these cases occur when the correlation increases from Phase I to Phase II, as illustrated in the left hand plots of Figures 8 and 9. In contrast to the Hotelling T^2 chart, the Verdier chart's *ARL*s are always significantly shorter than the target ARL_0. This property is desirable.

Second, the Hotelling T^2 chart seems to slightly outperform the Verdier chart in the cases, which preserve both the mean vector $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{\Sigma}$, so in particular the Pearson correlation.

While the Hotelling T^2 chart had shorter out-of-control estimates in nine cases, they are only on average about 14 *ARL*-steps shorter. Whereas the six times the Verdier chart outperformed the Hotelling T^2 chart, it is on average about 240 *ARL*-steps shorter. In the twelve cases where the Hotelling T^2 does not overshoot the in-control average run-length, the Verdier chart is on average approximately 80 *ARL*-steps shorter. Overall, the *Verdier* chart performs *better*.

TABLE 3 Weighted Vuong scheme *gaussian*(0.3) to *gaussian*(0.6) with confidence level 99%

P I Cop.	P II Cop.	λ_{ES}	α_v	ARL_0	ARL
gss(0.3)	*gss*(0.6)	NA	0.000275	365.31	136.19 ∈ [132.74; 139.65]
gss(0.3)	*gss*(0.6)	0.1	0.000095	366.95	112.23 ∈ [109.39; 115.09]
gss(0.3)	*gss*(0.6)	0.2	0.000143	367.27	131.22 ∈ [127.92; 134.52]
gss(0.3)	*gss*(0.6)	0.3	0.000107	366.03	119.86 ∈ [116.87; 122.85]
gss(0.3)	*gss*(0.6)	0.4	0.000107	363.02	114.81 ∈ [112.00; 117.63]
gss(0.3)	*gss*(0.6)	0.5	0.000137	363.95	**104.99** ∈ [102.39; 107.60]

The bold values signify the most advantageous ARLs. In this case the smallest ones.

TABLE 4 Weighted Vuong scheme *gaussian*(0.6) to *gaussian*(0.9) with confidence level 99%

P I Cop.	P II Cop.	λ_{ES}	α_v	ARL_0	ARL
gss(0.6)	*gss*(0.9)	NA	0.000130	363.32	26.76 ∈ [26.21; 27.30]
gss(0.6)	*gss*(0.9)	0.1	0.000069	363.85	**24.04** ∈ [23.56; 24.52]
gss(0.6)	*gss*(0.9)	0.2	0.000081	366.79	27.02 ∈ [26.46; 27.58]
gss(0.6)	*gss*(0.9)	0.3	0.000079	365.88	26.50 ∈ [25.98; 27.02]
gss(0.6)	*gss*(0.9)	0.4	0.000084	363.37	25.24 ∈ [24.75; 25.73]
gss(0.6)	*gss*(0.9)	0.5	0.000099	366.92	**23.90** ∈ [23.47; 24.33]

The bold values signify the most advantageous ARLs. In this case the smallest ones.

TABLE 5 Weighted Vuong scheme *gaussian*(0.9) to *gaussian*(0.6) with confidence level 99%

P I Cop.	P II Cop.	λ	α_v	ARL_0	ARL
gss(0.9)	*gss*(0.6)	NA	0.000816	364.51	**797.61** ∈ [776.55; 818.67]
gss(0.9)	*gss*(0.6)	0.1	0.000473	367.49	1090.59 ∈ [1062.72; 1118.46]
gss(0.9)	*gss*(0.6)	0.2	0.000671	363.75	907.81 ∈ [884.39; 931.24]
gss(0.9)	*gss*(0.6)	0.3	0.000641	366.20	886.29 ∈ [863.28; 909.29]
gss(0.9)	*gss*(0.6)	0.4	0.000656	363.58	860.23 ∈ [837.95; 882.52]
gss(0.9)	*gss*(0.6)	0.5	0.000671	366.34	844.51 ∈ [823.05; 865.97]

The bold values signify the most advantageous ARLs. In this case the smallest ones.

6.2 | Weighted copula estimation and the Vuong test

The weighted Vuong scheme described in Section 4.4 shows that weighting has a significant influence on the out-of-control ARL, confer Tables 3–5 as well as Figure 10. Overall the out-of-control performance of the weighted Vuong scheme is very sensitive to the type of shift in the correlation. For upward shifts, very short out-of-control ARLs are achieved, see Tables 3 and 4. However, for downward shifts it exhibits high inertia and signals very late, after more than twice the calibrated in-control time, see Table 5.

Three shortcomings present themselves

(1) Calibrating this model in the Vuong type I error parameter α_v is very time-consuming and resource intensive. The analysis of the Vuong approach took a total of 14,000 CPU hours, roughly 8500 of which are needed for calibration in the α_v Type-I-error parameter of the Vuong test statistic to achieve a calibrated false-alarm rate of 365 ± 2.5 corresponding to false-alarm rate $\alpha \in [0.00272; 0.00276]$. The out-of-control average run-length estimations roughly took 5500 CPU hours.
(2) In the bivariate normal case of decreasing the correlation from 0.9 to 0.6, this scheme takes at least twice as long to detect the out-of-control distribution compared to the in-control average run-lengths, as illustrated by Table 5. This is highly undesirable, especially considering the previous results for the Verdier chart.
(3) In nine out of the 15 design matrix entries pertaining to a change in the dependence structure, the out-of-control *ARL* of the *unweighted* scheme is notably smaller than any of the weighted cases. One of the nine cases is exemplified in Table 5 for all weighting parameters $\lambda \in \{\text{NA}, 0.1, 0.2, 0.3, 0.4, 0.5\}$. Furthermore, the weighted cases that are

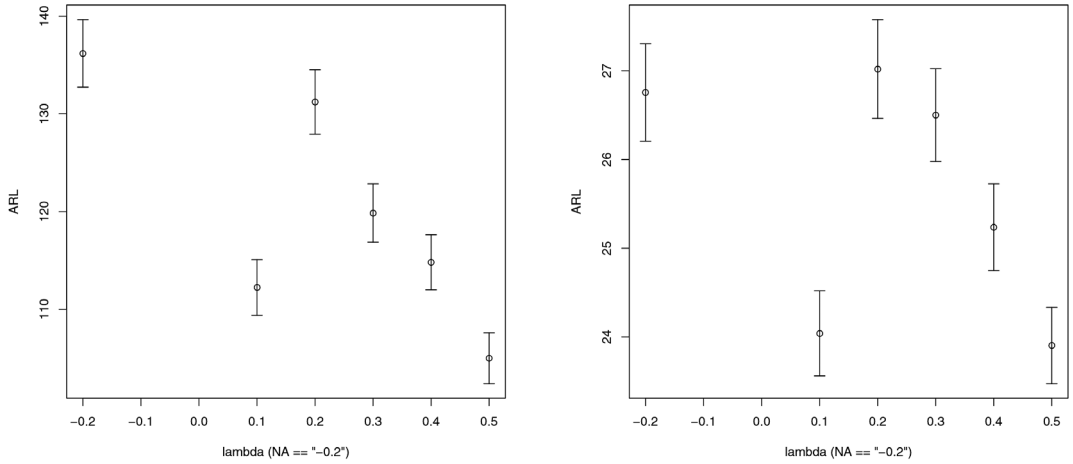

$gss(0.3)$ to $gss(0.6)$ $gss(0.6)$ to $gss(0.9)$

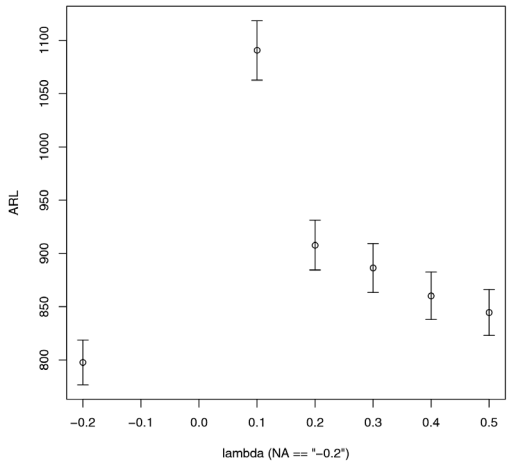

$gss(0.9)$ to $gss(0.6)$

FIGURE 10 Weighted Vuong scheme: confidence intervals for level 99%. The unweighted case $\lambda = \text{NA}$ is encoded as -0.2

significantly better than the unweighted cases did not seem to follow any discernible pattern. Sometimes smaller weights are better, sometimes larger ones, sometimes weights in-between.

Considering the discussion of the Vuong test in Section 4.4, the problems (2) and (3) have to be assumed to stem from the small sample size. Seemingly, the results from the asymptotic cannot be applied to small sample sizes in a straightforward manner. Lower sample size bounds for the validity of Vuong's test scheme have not been considered in the literature and are an important subject matter for future research.

The combination of weighted copula estimation and the Vuong approach the sample size of $N = 10$ seems to be far too small for bivariate copula estimation to be useful in the context of such severely constrained sample sizes, even with a very strong assumption of knowing both the Phase I distribution and the Phase II marginals.

7 | CONCLUSION AND OUTLOOK

We have contrasted the Hotelling T^2 chart with the monitoring chart proposed by Verdier[10] using an ample variety of out-of-control situations reflecting different copula-based association structures. It turns out that the Verdier chart is a robust solution with respect to the out-of-control *ARL* for a variety of out-of-control cases whereas T^2 exhibits excessively large *ARL*s in several out-of-control situations. We have demonstrated that the Vuong[14] test in combination with weights requires further investigation with respect to larger sample sizes and the type of shifts it detects well.

Our results demonstrate that copula-based reasoning is fruitful for the analysis of design of process monitoring techniques. A variety of issues remains for future research

(i) the analysis of Phase I estimation errors with respect to their influence on the in-control and out-of-control *ARL*s,
(ii) the extension to higher and possibly large dimensions, as well as
(iii) methods of signal tracking.

ACKNOWLEDGEMENTS

We gratefully acknowledge the compute resources and support provided by the University of Würzburg IT Centre and the German Research Foundation (DFG) through Grant No. INST 93/878-1 FUGG.

DATA AVAILABILITY STATEMENT

The data that support the introductory stock price example of this paper are available from Thomson Reuters.[11] Restrictions apply to the availability of these data, which were used under license for this study. Further details concerning availability are in Reference section.[11]

The data that support the findings concerning Verdier's monitoring chart and this paper's weighted Vuong approach in Section 6 are available from the corresponding author upon reasonable request.

ORCID

Andrew Easton https://orcid.org/0000-0002-6001-5760
Okki van Dalen https://orcid.org/0000-0002-3885-8155
Rainer Goeb https://orcid.org/0000-0003-4596-4068

Alessandro Di Bucchianico https://orcid.org/0000-0003-1842-5245

REFERENCES

1. Cherubini U, Luciano E, Vecchiato W. *Copula Methods in Finance*. Wiley; 2014.
2. Schoelzel C, Friederichs P. Multivariate non-normally distributed random variables in climate research – introduction to the copula approach. *Nonlin Process Geophys.* 2008;15(5):761–772. https://hal-cea.archives-ouvertes.fr/cea-00440431.
3. Ren X, Tian Y, Li S. Vine copula-based dependence description for multivariate multimode process monitoring. *Ind Eng Chem Res.* 2015;54(41):10001–10019. http://doi.org/10.1021/acs.iecr.5b01267.
4. Czado C. *Analyzing Dependent Data with Vine Copulas*. Vol 222. Springer; 2019: https://doi.org/10.1007/978-3-030-13785-4.
5. Busababodhin P, Amphanthong P. Copula modelling for multivariate statistical process control: a review. *Commun Stat Appl Methods.* 2016;23:497–515. http://www.csam.or.kr/journal/view.html?doi=10.5351/CSAM.2016.23.6.497.
6. Hryniewicz O, Szediw A. Sequential signals on a control chart based on nonparametric statistical tests. In: *Frontiers in Statistical Quality Control*. 2010;9:99–117.
7. Hryniewicz O. On the robustness of the shewhart control chart to different types of dependencies in data. In: *Frontiers in Statistical Quality Control*. 201210:19–33.
8. Xie Y, Xie M, Goh TN. Two MEWMA charts for Gumbel's bivariate exponential distribution. *J Qual Technol*. 2011;43:50–56.
9. Kuvattana S, Sukparungsee S, Areepong Y, Busababodhin P. Bivariate copulas on the exponentially weighted moving average control chart. *Songklanakarin J Sci Technol.* 2016;38:569–574.
10. Verdier G. Application of copulas to multivariate control charts. *J Stat Plan Infer*. 2013;143(12):2151–2159.
11. Thomson Reuters Eikon. [dax-30 daily unadjusted prices 2014-11-28 through 2019-11-28.]. 2019. https://eikon.thomsonreuters.com/index.html. Accessed December 12, 2019.
12. Sklar A. *Fonctions de Répartition à n Dimensions et Leurs Marges*. Vol 8. Publications de Institut de Statistique de Université de Paris; 1959:229–231.
13. Nelsen R. *An Introduction to Copulas*. Springer Science & Business Media; 2007.
14. Vuong Q. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*. 1989;57(2):307–333.
15. Mason RL, Young JC. *Multivariate Statistical Process Control with Industrial Applications*. Philadelphia, PA: Americal Statistical Association – Society for Industrial and Applied Mathematics, SIAM; 2002:19104. http://doi.org/10.1137/1.9780898718461.
16. Baíllo A, Cuevas A. Parametric versus nonparametric tolerance regions in detection problems. *Comput Stat*. 2006;21(3–4):523–536.
17. Cadre B, Pelletier B, Pudlo P. Estimation of density level sets with a given probability content. *J Nonparametric Stat*. 2013;25(1):261–272.
18. Mühlig B. *Multivariate process monitoring based on copula structures* [master's thesis]. Department of Statistics, Julius-Maximilians-Universität Würzburg; 2017.
19. Nagler T, Vatter T. rvinecopulib: high performance algorithms for vine copula modeling. R package version 0.5.3.1.0. 2020. https://cran.r-project.org/package=rvinecopulib.

20. Brechmann EC, Czado C, Aas K. Truncated regular vines and their applications. *Can J Stat*. 2012;40(1):68–85.
21. Schepsmeier U, Nagler T, Stoeber J, Brechmann EC, Graeler B, Erhardt T, Almeida C, Min A, Czado C, Hofmann M, Killiches M, Joe H, Vatter T. (2019) VineCopula: statistical inference of vine copulas. R package version 2.3.0. https://cran.r-project.org/package=VineCopula.
22. Amisano G, Giacomini R. Comparing density forecasts via weighted likelihood ratio tests. *J Bus Econ Stat*. 2007;25(2):177–190.

AUTHOR BIOGRAPHIES

Andrew Easton holds an M.Sc. degree in Mathematics from the University of Würzburg, Germany. His research interest is the application of copula theory to process monitoring in financial and industrial contexts with the goal of refining dependency models. Currently he is pursuing his research interests after gaining a year of practial experience with regards to statistical analysis of health insurance claims at HUK-COBURG.

Okki van Dalen holds an M.Sc. degree in Industrial and Applied Mathematics from Eindhoven University of Technology. His main area of interest is the application of machine learning and deep learning algorithms in order to gain valuable insights into both industrial and business processes, especially statistical process control and predictive maintenance. Currently he is a Data Scientist at BAS Trucks.

Rainer Goeb is a Professor of Statistics at the Institute of Matematics of the University of Würzburg, Germany. His research interest is in the industrial application of statistics, in particular statistical sampling, process monitoring, predictive analytics, risk analysis and the standardisation of statistical methods. He is the chair of the subcommittee "Acceptance Sampling" at the Technical Committee 69 "Application of Statistical Methods" of the International Organization for Standardization.

Alessandro Di Bucchianico holds an M.Sc. degree in mathematics from Amsterdam University and a Ph.D. degree in mathematics from Groningen University, the Netherlands. Currently he is an Associate Professor of Statistics at the Mathematics and Computing Science Department of the Eindhoven University of Technology. His main research interests are statistical process control, reliability (both hardware and software reliability) and teaching of statistics.

How to cite this article: Easton A, van Dalen O, Goeb R, Di Bucchianico A. Online monitoring of process variance drifts for leptokurtic and heavy-tailed data. *Qual Reliab Eng Int*. 2022;38:1272–1288. https://doi.org/10.1002/qre.3034