Change-point analysis for dependence structures in finance and insurance*

Alexandra Dias(1) and Paul Embrechts(2)

Department of Mathematics
ETH Zürich
CH–8092 Zürich
Switzerland

Abstract

Over the recent years, both in finance and insurance, the modelling of dependence beyond linear correlation has become a key area of research. The notion of copula has been used with success in order to model these more general dependence concepts. We will discuss changes in dependence structures by using change-point techniques for specific parametric copula families. Besides some basic theory, some applied examples will be presented.

Keywords: copula, change–point, likelihood–ratio, bootstrap

AMS Classification: 62F03, 62F40, 62P05

*Paper based on talks given by the authors: (1) at the "Sixth International Congress on Insurance: Mathematics and Economics, Lisbon, 2002" and (2) at the "IX Congresso Anual da Sociedade Portuguesa de Estatística, Ponta Delgada, 2001." The first author would like to thank the support from Fundação para a Ciência e a Tecnologia - FCT/POCTI and Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Portugal. Also, support from the National Centre of Competence in Research "Financial Valuation and Risk Management" (NCCR FINRISK) is gratefully acknowledged.

1 Introduction

In risk management it is common to have to deal with multivariate risks and it is well known that the homogeneity of the marginal behaviour of each risk does not imply a global homogeneous behaviour. For example, there is considerable interest in the dynamic behaviour of correlation between different risks as a function of time; see for instance Boyer et al. (1999), Loretan and Phillips (1994) and Longin and Solnik (2001). Because of the fundamental importance of the notion of linear correlation in finance and insurance, such changes may have a non-trivial impact on the pricing and hedging of underlying instruments, or of the risk measurement thereof. As a consequence, a more systematic study for the dynamic behaviour of the dependence structure underlying multivariate risks is called for. Our paper will concentrate on the notion of copula as for instance discussed in Embrechts et al. (2002), Embrechts et al. (2000) and the references therein. For a specific application of copula modelling to multiline insurance products, see Blum et al. (2002). The reader is referred to these papers for the notation, definitions and basic results. In our paper we will concentrate on parametric copula modelling of dependence and discuss change–point analysis questions in this framework. For some related work, see for instant Gombay and Horváth (1999). For a detailed treatment of the change–point theory underlying our approach, see Csörgő and Horváth (1997) and references therein.

The paper is organised as follows: in Section 2 we summarise the maximum likelihood approach of change–point analysis with the detection of changes in the dependence structure in mind. We compute the distribution of the resulting test statistic for some commonly used parametric copulas, infer about the power of the test, discuss the construction of confidence intervals and illustrate the methods introduced on a simulated data example. Section 3 has an analysis for two financial positions on the same simulated data, using the methodologies presented in the previous section. An application of the introduced methods on real data is given in Section 4. We conclude with a brief summary in Section 5.

2 Statistical change-point analysis

2.1 The test statistic

Suppose that we have n vectors of observations, each composed by m risks. Formally, let $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$ be a sequence of independent random vectors in $\mathbb{R}^m$ with distribution functions $F(\mathbf{x}; \boldsymbol{\theta}_1, \boldsymbol{\eta}_1), \ldots, F(\mathbf{x}; \boldsymbol{\theta}_n, \boldsymbol{\eta}_n)$, respectively, where $\boldsymbol{\theta}_i$ and $\boldsymbol{\eta}_i$ are parameters of the distribution functions such that $\boldsymbol{\theta}_i \in \Theta^{(1)} \subseteq \mathbb{R}^d$ and $\boldsymbol{\eta}_i \in \Theta^{(2)} \subseteq \mathbb{R}^p$ for $1 \leq i \leq n$. We will be primarily interested in a change–point analysis for the θ_i's, whereas the η_i's will be nuisance parameters. As a consequence, we are interested in testing the null hypothesis

$$H_0 : \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2 = \ldots = \boldsymbol{\theta}_n \qquad \text{and} \qquad \boldsymbol{\eta}_1 = \boldsymbol{\eta}_2 = \ldots = \boldsymbol{\eta}_n$$

versus the alternative

$$H_A : \boldsymbol{\theta}_1 = \ldots = \boldsymbol{\theta}_{k^*} \neq \boldsymbol{\theta}_{k^*+1} = \ldots = \boldsymbol{\theta}_n \quad \text{and} \quad \boldsymbol{\eta}_1 = \boldsymbol{\eta}_2 = \ldots = \boldsymbol{\eta}_n.$$

Here k^* is the location or time of the (assumed) single change–point. All the parameters $(\boldsymbol{\theta}, \boldsymbol{\eta}) \in \Theta^{(1)} \times \Theta^{(2)}$ are supposed to be unknown under both hypotheses. As a start, assume that $k^* = k$ is known. In that case, the question consists of testing if two samples come from the same population and can be done through the generalised likelihood ratio test. The null hypothesis will be rejected for small values of the test statistic

$$\Lambda_k = \frac{\sup\limits_{(\boldsymbol{\theta},\boldsymbol{\eta}) \in \Theta^{(1)} \times \Theta^{(2)}} \prod\limits_{1 \leq i \leq n} f(\mathbf{X}_i; \boldsymbol{\theta}, \boldsymbol{\eta})}{\sup\limits_{(\boldsymbol{\theta},\boldsymbol{\theta}',\boldsymbol{\eta}) \in \Theta^{(1)} \times \Theta^{(1)} \times \Theta^{(2)}} \prod\limits_{1 \leq i \leq k} f(\mathbf{X}_i; \boldsymbol{\theta}, \boldsymbol{\eta}) \prod\limits_{k < i \leq n} f(\mathbf{X}_i; \boldsymbol{\theta}', \boldsymbol{\eta})}.$$

As the estimation of Λ_k is carried out through maximum likelihood, all the necessary conditions of regularity and efficiency have to be assumed (see for instance Lehmann (1991)).

If we denote

$$L_k(\boldsymbol{\theta}, \boldsymbol{\eta}) = \sum_{1 \leq i \leq k} \log f(\mathbf{X}_i; \boldsymbol{\theta}, \boldsymbol{\eta})$$

and

$$L_k^*(\boldsymbol{\theta}, \boldsymbol{\eta}) = \sum_{k < i \leq n} \log f(\mathbf{X}_i; \boldsymbol{\theta}, \boldsymbol{\eta})$$

then the likelihood ratio equation can be written as

$$-2\log(\Lambda_k) = 2[L_k(\hat{\boldsymbol{\theta}}_k, \hat{\boldsymbol{\eta}}_k) + L_k^*(\boldsymbol{\theta}_k^*, \hat{\boldsymbol{\eta}}_k) - L_k(\hat{\boldsymbol{\theta}}_n, \hat{\boldsymbol{\eta}}_n)].$$

As stated, the statistic Λ_k tests H_0 for k known. In the (more realistic) case when k is unknown, H_0 will be rejected for large values of

$$Z_n = \max_{1\leq k<n}(-2\log(\Lambda_k)). \tag{1}$$

The asymptotic distribution of $Z_n^{1/2}$ can be derived using Extreme Value Theory (EVT) techniques; see Embrechts et al. (1997) and Embrechts (2000). Indeed, let

$$A(x) = (2\log(x))^{1/2}$$

and

$$D_d(x) = 2\log(x) + \frac{d}{2}\log(\log(x)) - \log(\Gamma(d/2)),$$

where $\Gamma(t)$ is the usual gamma function

$$\Gamma(t) = \int_0^\infty y^{t-1}\exp(-y)\,dy.$$

Then, if H_0 and all the necessary regularity conditions hold, we have that

$$\lim_{n\to\infty} P(A(\log(n))Z_n^{1/2} \leq t + D_d(\log(n))) = \exp(-2\exp(-t)) \tag{2}$$

for all t. The right side of (2) is the square of a Gumbel distribution function. It is known that the rate of convergence in results like (2) is very slow, see for instance Embrechts et al. (1997), page 150, and therefore care has to be taken to base a hypothesis test on the asymptotic distribution, especially for small and moderate sample sizes. For this reason Gombay and Horváth (1996) derived a result on which a test can be based that yields better rejection regions for smaller sample sizes. Under H_0 and supposing that all the necessary regularity conditions hold, for $n\to\infty$,

$$\left| Z_n^{1/2} - \sup_{1/n\leq t\leq 1-1/n}\left(\frac{B_n^{(d)}(t)}{t(1-t)}\right)^{1/2}\right| = o_P(\exp(-(\log(n))^{1-\varepsilon})) \tag{3}$$

for all $0<\varepsilon<1$, where $\{B_n^{(d)} : 0\leq t\leq 1\}$ is a sequence of stochastic processes such that $\{B_n^{(d)} : 0\leq t\leq 1\} \overset{d}{=} \{B^{(d)} : 0\leq t\leq 1\}$ for each n

and $B^{(d)}(t)=\sum_{1\le i\le d} B_i^2(t)$, where $\{B_s(t):0\le t\le 1\}$, $s=1,\ldots,d$ are independent Brownian bridges.

For $0<\alpha<1$ let

$$z_n=z_n(1-\alpha)=\sup\left\{x\ge 0: P(Z_n^{1/2}\le x)\le 1-\alpha\right\}$$

and

$$u(h,l)=u(h,l;1-\alpha)=\sup\left\{x\ge 0: P\left(\sup_{h\le t\le 1-l}\left\{\frac{B^{(d)}(t)}{t(1-t)}\right\}^{1/2}\le x\right)\le 1-\alpha\right\}.$$

As a consequence of (2) and (3), if $h(n)\ge 1/n$, $l(n)\ge 1/n$ and for $n\to\infty$,

$$(2\log\log n)^{-1/2}\sup_{1-c(n)\le t\le 1-/n}\left\{\frac{B_n^{(d)}(t)}{t(1-t)}\right\}\xrightarrow{P}1-\varepsilon^*$$

is satisfied for some $0<\varepsilon^*\le 1$, where

$$c(n)=\exp((\log n)^{1-\varepsilon^*})/n,$$

we have that

$$\lim_{n\to\infty}P(Z_n^{1/2}>u(h(n),l(n)))=\alpha$$

and

$$|z_n(1-\alpha)-u(h(n),l(n))|=o((\log\log n)^{-1/2});$$

see Csörgő and Horváth (1997), page 24. The distribution function of $\sup_{h\le t\le l}\{B^{(d)}(t)/(t(1-t))\}^{1/2}$ has no simple closed form. For practical applications we use, for $0<h<l<1$, the following approximation:

$$P\left(\sup_{h\le t\le 1-l}\left\{\frac{B^{(d)}(t)}{t(1-t)}\right\}^{1/2}\ge x\right)=\frac{x^d\exp(-x^2/2)}{2^{d/2}\Gamma(d/2)}\times$$
$$\left(\log\frac{(1-h)(1-l)}{hl}-\frac{d}{x^2}\log\frac{(1-h)(1-l)}{hl}+\frac{4}{x^2}+\mathcal{O}\left(\frac{1}{x^4}\right)\right),\quad (4)$$

as $x\to\infty$; see Gombay and Horváth (1996).

Based on (2) and (3) asymptotic critical values can be computed. Table 1 has the values computed from (2) in column $z^{(1)}$ and from (3) in column $z^{(2)}$. Gombay and Horváth (1996) found that $h(n)=l(n)=(\log n)^{3/2}/n$ makes $u(h,l)$ a good approximation for $z_n=z_n(1-\alpha)$. We use the same choice to obtain the asymptotic critical values $z^{(2)}$ listed in Table 1. On the other hand, for a given model, we can perform Monte Carlo simulations in order to get a further approximation to the critical values of the likelihood ratio test statistic $Z_n^{1/2}$. Under the null hypothesis, $5'000$ repetitions of $Z_n^{1/2}$ for $n = 50, 100$ and 500 were performed in the case of bivariate Gumbel, Frank and Gaussian copulas; see Joe (1997), Nelsen (1999) or Embrechts et al. (2002) for their definitions and the parameterisations used. If we denote the simulated repetitions of $Z_n^{1/2}$ by $Z_{n,t}^{1/2^*}$, $t = 1, \ldots, N$, then the estimated critical value z_n^* at the level $1 - \alpha$ is given by $Z_{n,((1-\alpha)(N+1))}^{1/2^*}$ where $Z_{n,(r)}^{1/2^*}$ is the rth ordered value. (For an explanation on Monte Carlo quantiles, see for instance Hall (1992)) These results are also to be found in Table 1.

Sample size	$1-\alpha$	$z^{(1)}$	$z^{(2)}$	z_n^{Gu}	z_n^{Fr}	z_n^{Ga}
	0.90	3.18	2.69	2.67	2.51	2.87
50	0.95	3.62	2.97	2.93	2.76	3.13
	0.99	4.60	3.52	3.45	3.22	3.59
	0.90	3.23	2.79	2.81	2.64	2.95
100	0.95	3.64	3.06	3.02	2.88	3.19
	0.99	4.57	3.59	3.54	3.39	3.78
	0.90	3.31	2.95	2.94	2.89	3.07
500	0.95	3.69	3.20	3.18	3.13	3.32
	0.99	4.54	3.71	3.61	3.68	3.79

Table 1: *Critical values for the likelihood ratio test $Z_n^{1/2}$ given by (1). The third column has the asymptotic critical values given by (2), in the fourth column the $z^{(2)}$'s are derived from (3). The z_n^{Gu}, z_n^{Fr} and z_n^{Ga} are the simulated critical values for the Gumbel, Frank and Gaussian bivariate copulas, respectively. The simulations are made under H_0 and one parameter can change under the alternative hypothesis ($d = 1$).*

Table 1 can be used to test whether the parameter of the Gumbel or the Frank copulas changed for a given data set. Or we can test if the correlation in a Gaussian copula is constant. In Csörgő and Horváth (1997) the asymptotic values $z^{(1)}$ and $z^{(2)}$ are compared with univariate normal, exponential and Poisson observations. In Table 1, we compare the same asymptotic critical values with those obtained from bivariate Gumbel, Frank and Gaussian copula observations. The simulations show that in fact the true critical values are smaller than the ones derived from (2) meaning that the later result usually provides a conservative rejection region. In the case of univariate normal, exponential and Poisson observations, this fact was already observed by Csörgő and Horváth (1997). In Figure 1 we can see that,

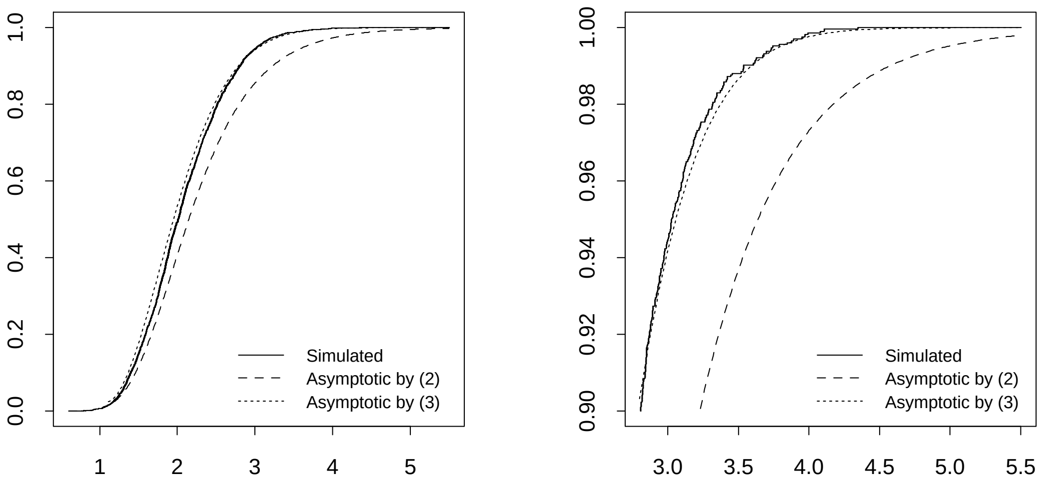

Figure 1: *The distribution function of* $Z_{100}^{1/2}$ *given by the asymptotic approximations (2) and (3) and by simulation for the Gumbel case. The left panel presents the full distribution functions whereas the right panel concentrates on the tail region above 90%.*

for the Gumbel case with sample size 100, the asymptotic result (3) is much closer to the distribution function of the test statistic than (2). In the right panel are the values of the same distribution functions for more than 0.9 of probability. The asymptotic result (3) still gives a good approximation in the tail. For the three sample sizes in Table 1 and for the Gumbel and the Frank distributions, we can observe in Figure 2 the asymptotic and the simulated distribution functions of $Z_n^{1/2}$. In these plots we see that the asymptotic result (3) actually gives a good (small sample) approximation for the rejection regions and is much better than the approximation given by (2), for all cases considered.

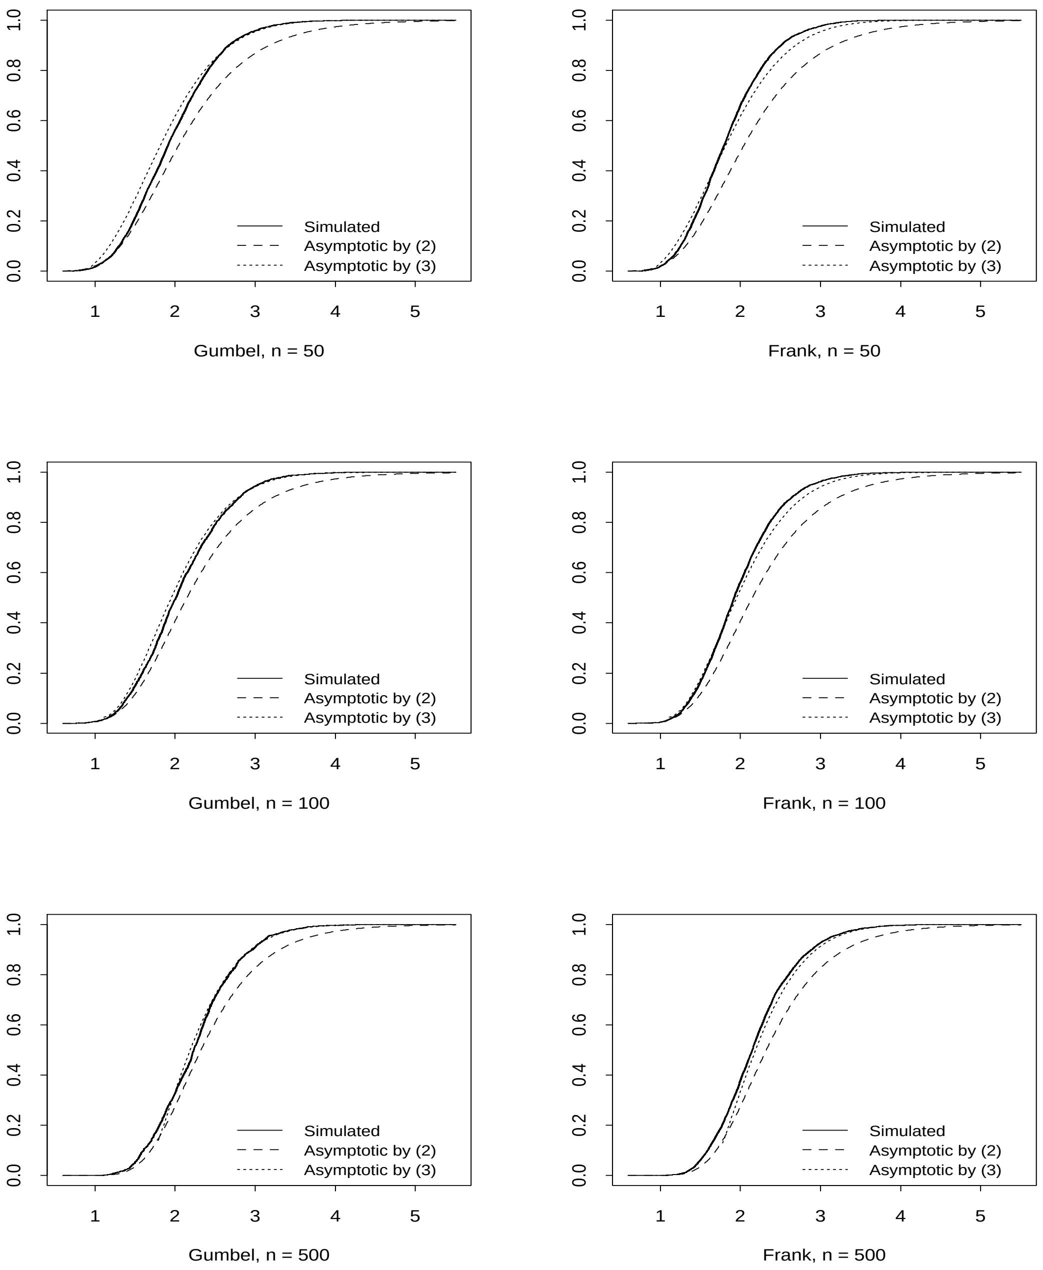

Figure 2: *The distribution function of* $Z_n^{1/2}$ *given by the asymptotic results (2) and (3) and by simulation. The left panels correspond to the Gumbel distribution whereas the right ones to the Frank. The rows of the graphics correspond to different sample sizes* $n = 50$, $n = 100$ *and* $n = 500$.

2.2 An example: the Gumbel case

A set of 500 observations was simulated from a bivariate Gumbel copula with parameter $\theta = 0.8$ for the first 250 observations and $\theta = 0.4$ for the second half. Recall from Embrechts et al. (2002) that a random vector (U, V) has a Gumbel copula if for $0 \leq u, v \leq 1$,

$$P(U \leq u, V \leq v) = \exp\left\{ - \left[(-\ln u)^{1/\theta} + (-\ln v)^{1/\theta} \right]^{\theta} \right\}, \quad \theta \in (0, 1].$$

In the case of the Gumbel copula, there is a straightforward link between the parameter θ and the Kendall–τ correlation coefficient,

$$\tau = 1 - \theta,$$

hence this simulation can also be interpreted as testing for a change in τ from 0.2 to 0.6. Questions of the latter type are important in various applications, see for instance Gombay and Horváth (1999). We can at first assume that all the data come from a homogeneous model, following a Gumbel copula with unknown parameter θ, say. Fitting this model by maximum likelihood produced an estimate for the copula parameter of $\hat{\theta} = 0.6325$. Figure 3 has the scatterplot of the simulated data and a qq–plot of the fitted against the theoretical quantiles. The qq-plot points to a fairly good fit, as the pairs of quantiles lay almost all very close to the main diagonal. Actually, we can test whether $V|U$ has indeed a standard uniform distribution. A Kolmogorov–Smirnov goodness of fit test does not reject the fitted Gumbel model (p–value of 0.22). This points to an important property of copula modelling: through a Gumbel model for instance, a whole range of linear correlation values in agreement with the data can be covered. We will refrain from discussing this point here any further but will return to it in future publications.

We now turn to the main issue underlying the example, the performance of the test statistic Λ_k towards the change–point detection from $\theta = 0.8$ to 0.4 around the middle of the data set. Figure 4 is the plot of $(-2\log(\Lambda_k))^{1/2}$ for $1 \leq k < n$. The maximum of these values is $Z_{500}^{1/2} = 10.02$, reached at $k = 248$. Either by approximations (2), (3) or by simulation, $P(Z_{500}^{1/2} \leq 10.02) = 1$, hence we reject H_0 very significantly. Hence our analysis yields a change–point at observation $k = 248$ with pre– and post–values of θ estimated as $\hat{\theta}_b = 0.8515$, $\hat{\theta}_a = 0.4087$, respectively.

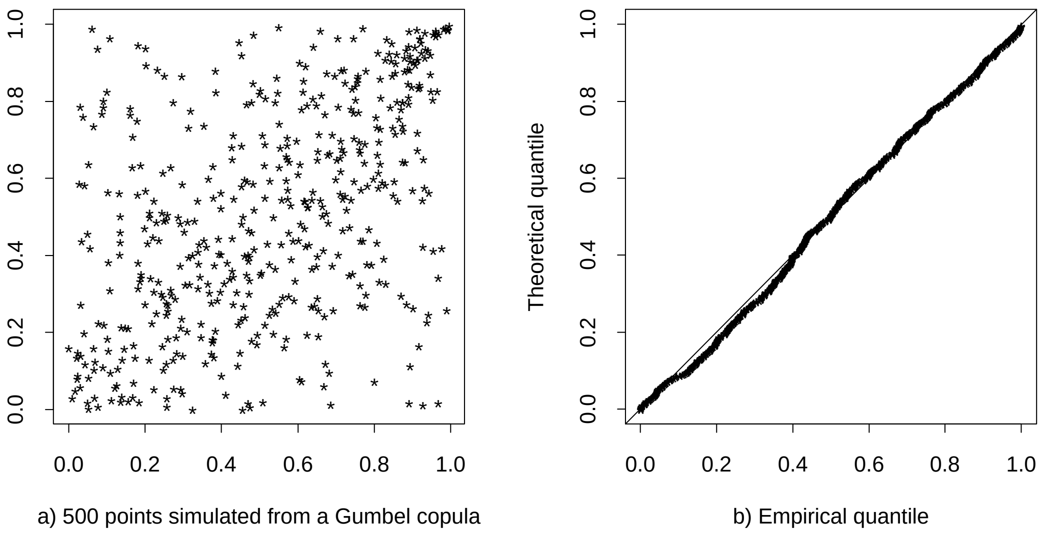

Figure 3: *Simulated values from a Gumbel copula with the θ–values 0.8 and 0.4 and qq–plot corresponding to the maximum likelihood fit, where* $\hat{\theta} = 0.6325$.

The excellent performance of the test in detecting the simulated change in the example above has to be taken with caution. The analysis based on the assumption of homogeneity (no θ–change) already yields a warning sign. It turns out that the power function for $Z_n^{1/2}$ depends on the location of the change and on the size of the change. For instance, for a small change near the limits of the sampling time interval, the test doesn't give such good results. The next section is devoted to this issue. Furthermore, in most practical applications, interdependence between the data may further weaken the analysis.

2.3 The power of the test

As already discussed above in the Gumbel example, it will be important to analyse the power of the change–point test presented. The power function is defined as follows $\beta(\alpha) = P(Z_n^{1/2} > z_n(\alpha)|H_A)$ where H_A stands for the alternative hypothesis of one change–point. Hence we need the distribution function of the test statistic $Z_n^{1/2}$ under the alternative hypothesis; in the case of interest, i.e. the bivariate Gumbel and Frank copulas, this distribution

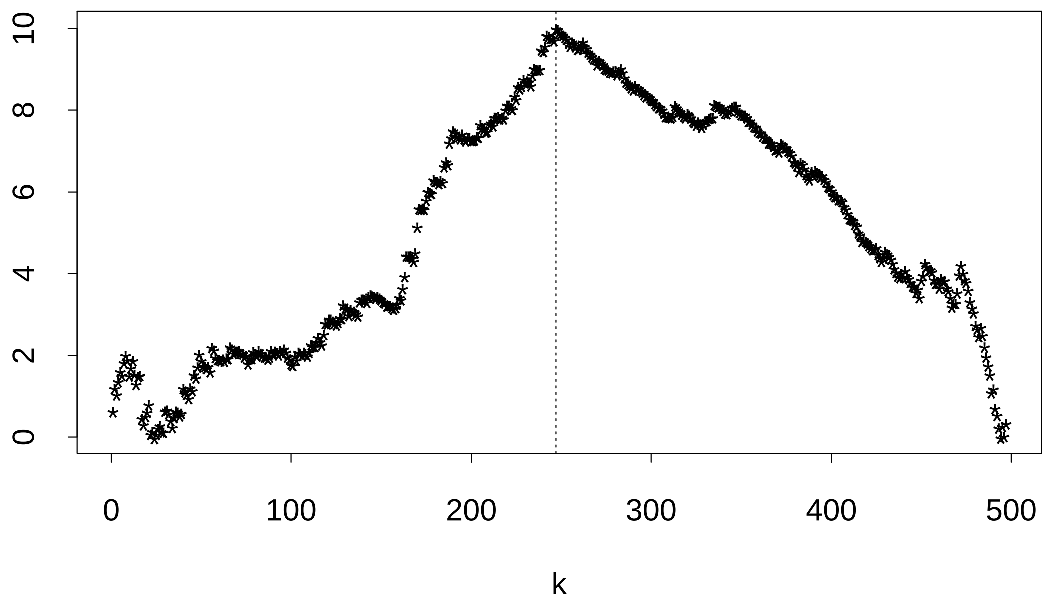

Figure 4: *Values of* $(-2\log(\Lambda_k))^{1/2}$, $1 \leq k < 500$, *for the simulated data in the Gumbel example of Section 2.2.*

is unknown. We therefore perform Monte Carlo simulations to estimate the power function and replicate $Z_n^{1/2}$, under the alternative, $N = 5'000$ times for the Gumbel copula and a sample size of $n = 100$. The power function, for a given level $1-\alpha$, is estimated by $\hat{\beta}_n(\alpha) = (1+\#\{Z_n^{1/2^*} > z_n(\alpha)\})/(1+N)$. As we know that the distribution function of $Z_n^{1/2}$ depends also on the size of the change and on the location of the change, under the alternative hypothesis, we simulated several different scenarios. For the location of the change k^*, we took values $10, 25$ and 50. As for the size of the change, we considered the Gumbel distribution with parameter $0 < \theta \leq 1$ and changes between 0.1 and 0.9. The results reported in Table 2 are based on the critical values given in Table 1 which were obtained by the asymptotic distribution in (3).

k^*	Level $1-\alpha$	Size of the change 0.1	0.3	0.5	0.9
	0.90	0.1330	0.4028	0.8416	1.000
10	0.95	0.0618	0.2642	0.7341	1.000
	0.99	0.0130	0.0890	0.4423	1.000
	0.90	0.1748	0.7314	0.9984	1.000
25	0.95	0.0932	0.6092	0.9928	1.000
	0.99	0.0202	0.3466	0.9596	1.000
	0.90	0.2030	0.8650	0.9998	1.000
50	0.95	0.1140	0.7902	0.9994	1.000
	0.99	0.0290	0.5846	0.9970	1.000

Table 2: *Power function values of the change–point test for different locations k^* and sizes of the change. The sample size is $n = 100$. Each value was obtained from $5'000$ simulations of the corresponding case.*

The values in Table 2 confirm that the power of the test strongly depends on the size of the change and on its location. As is to be expected, the bigger the size of the change, the more powerful the test is. The closer the change–point k^* comes to the edge of the data, the less powerful the test becomes. Based on Table 2, it seems that the loss of power is more due to a size change than to the location of the change. For applications this is relevant as it is more important to get an early warning once a change has taken place.

So far we have tested for the existence of a change–point. In the next section we will look more in detail at the estimation of the time of a change–point, also constructing confidence intervals for this time.

2.4 The time of the change and corresponding confidence intervals

If we assume that there is exactly one change–point, then the maximum likelihood estimator for the time of the change is given by

$$\hat{k}_n = \min\{1 \leq k < n : Z_n = -2\log(\Lambda_k)\}. \tag{5}$$

In the case that there is no change, $\hat{k}_n$ will take a value near the limits of the sample. This holds because under the null hypothesis and if all the necessary regularity conditions hold, for $n \to \infty$,

$$\hat{k}_n/n \xrightarrow{d} \xi, \tag{6}$$

where $P(\xi = 0) = P(\xi = 1) = 1/2$; see Csörgő and Horváth (1997), page 51. Figure 5 gives the frequency plot of $\max_{1 \le k < 500}(-2\log(\Lambda_k))$ for $5'000$ samples of size 500 simulated from a Gumbel copula, with parameter $\theta = 0.5$, under the null hypothesis of no–change. We can see from that graph that indeed the estimator of the time of the change mostly lays near the beginning or near the end of the time scale.

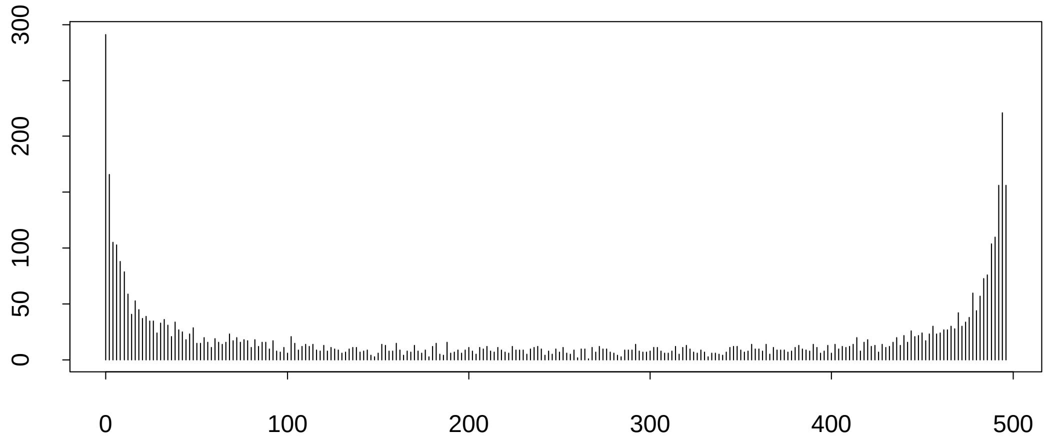

Figure 5: *Frequency plot of the time k for which* $Z_{500} = \max_{1 \le k < 500}(-2\log(\Lambda_k))$ *is attained, based on* $5'000$ *simulations of Gumbel copula samples of size n=500 with* $\theta = 0.5$ *(no change–point).*

The value $\hat{k}_n$ in (5) gives an estimate for the time of the change. In order to construct a confidence interval for the time of the change, we need to know or to approximate the distribution of $\hat{k}_n - k_0$, where k_0 is the true time of the change. One approximation can be obtained using bootstrap methodology. In the literature, there are several different approaches to construct bootstrap confidence intervals; see for instance Hall (1992), Efron and Tibshirani (1993) and Davison and Hinkley (1997). One of the simplest ways for constructing such an interval is the percentile method. A theoretical $(1 - 2\alpha)100\%$ confidence interval for the time of the change k_0 is of the form

$$(\hat{k}_n - a_2, \hat{k}_n - a_1) \tag{7}$$

where a_1 and a_2 satisfy

$$P(\hat{k}_n - a_2 \leq k_0 \leq \hat{k}_n - a_1) = 1 - 2\alpha.$$

Moreover if we want an equitailed interval, then we require that

$$P(k_0 \leq \hat{k}_n - a_2) = \alpha = P(k_0 > \hat{k}_n - a_1), \tag{8}$$

hence a_1 and a_2 are quantiles of the random variable $\hat{k}_n - k_0$.

Suppose that $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ is the sample from which we estimated the time of the change $\hat{k}_n$ assuming that the data came from a population with distribution function in the copula family C_θ. Then we have the maximum likelihood estimates $\hat{\theta}_b$ and $\hat{\theta}_a$ of the parameter of the distribution before and after the time of the change, respectively. From this we replicate the original sample simulating N samples of size n from the fitted distribution

$$F_{\hat{\theta}_b, \hat{\theta}_a, \hat{k}_n}(\mathbf{x}_i) = C_{\hat{\theta}_b}(\mathbf{x}_i)\mathbf{1}_{\{i \leq \hat{k}_n\}} + C_{\hat{\theta}_a}(\mathbf{x}_i)\mathbf{1}_{\{i > \hat{k}_n\}}$$

and compute the estimated time of the change $\hat{k}^*_{n,i}$ for each replicate sample, $i = 1, \ldots, N$. These replicates allow us to estimate the distribution function of $\hat{k}^*_n - \hat{k}_n$, where $\hat{k}^*_n$ is the time of the change estimate of a resample from a population with distribution function $F_{\hat{\theta}_b, \hat{\theta}_a, \hat{k}_n}$. Suppose that k^*_α and $k^*_{1-\alpha}$ are the quantiles of $\hat{k}^*_n$ such that

$$P(\hat{k}^*_n \leq k^*_\alpha) = \alpha = P(\hat{k}^*_n > k^*_{1-\alpha}),$$

then

$$P(k^*_\alpha - \hat{k}_n \leq \hat{k}^*_n - \hat{k}_n \leq k^*_{1-\alpha} - \hat{k}_n) = 1 - 2\alpha. \tag{9}$$

The bootstrap principle, for confidence intervals, consists on assuming that we can approximate the quantiles of $\hat{k}_n - k_0$ by the quantiles of $\hat{k}^*_n - \hat{k}_n$. Assuming that this bootstrap approximation works, from (8) and (9) we then have that

$$P(k^*_\alpha - \hat{k}_n \leq \hat{k}_n - k_0 \leq k^*_{1-\alpha} - \hat{k}_n) \approx 1 - 2\alpha,$$

or that

$$P(\hat{k}_n - (k^*_{1-\alpha} - \hat{k}_n) \leq k_0 \leq \hat{k}_n - (k^*_\alpha - \hat{k}_n)) \approx 1 - 2\alpha \tag{10}$$

where $(k^*_{1-\alpha} - \hat{k}_n)$ and $(k^*_\alpha - \hat{k}_n)$ are the bootstrap approximations of respectively a_2 and a_1 in (7) for the equitailed case.

It is well known that the estimation of confidence limits for a transformation of a parameter may give much better results than the direct estimation for the parameter; see for instance Efron and Tibshirani (1993), pages 54 and 162. Suppose that there exists a transformation of the random variable $\hat{k}_n$, say $\hat{u}_n = h(\hat{k}_n)$, which has a symmetric distribution. Then applying (10) we have that

$$P(\hat{u}_n - (u^*_{1-\alpha} - \hat{u}_n) \leq h(k_0) \leq \hat{u}_n - (u^*_\alpha - \hat{u}_n)) \approx 1 - 2\alpha, \qquad (11)$$

where u^*_α is the α–quantile of the transformed random variable $h(\hat{k}^*_n)$. Because of symmetry we know that $a_1 = -a_2$ in (7). So we can transform the approximated quantiles of $\hat{u}_n - h(k_0)$ in (11) such that

$$P(\hat{u}_n - (\hat{u}_n - u^*_\alpha) \leq h(k_0) \leq \hat{u}_n - (\hat{u}_n - u^*_{1-\alpha})) \approx 1 - 2\alpha.$$

Simplifying this expression and transforming back to the original scale we obtain

$$P(k^*_\alpha \leq k_0 \leq k^*_{1-\alpha}) \approx 1 - 2\alpha. \qquad (12)$$

In order to obtain the quantiles k^*_α from the bootstrap replicates $\hat{k}^*_{n,i}$ $(i = 1, \ldots, N)$ we note that, if $\hat{k}^*_{n,1}, \ldots, \hat{k}^*_{n,N}$ are independent and identically distributed with distribution function H, then

$$E[\hat{k}^*_{n,(j)}] = H^{-1}\left(\frac{j}{N+1}\right),$$

where, as before, $\hat{k}^*_{n,(j)}$ denotes the jth ordered value. Using this result an estimate for $k^*_\alpha = H^{-1}(\alpha)$ is $\hat{k}^*_{n,((N+1)\alpha)}$. Finally substituting these values in (12) we obtain that the bootstrap confidence interval for the time of the change k_0 becomes

$$\left(\hat{k}^*_{n,((N+1)\alpha)}, \hat{k}^*_{n,((N+1)(1-\alpha))}\right). \qquad (13)$$

The interval in (13) is often referred to as the bootstrap percentile confidence interval. Below, we highlight this construction on the example presented in Section 2.2.

Example: In our simulated example, the estimate for the time of the change is $\hat{k}_{500} = 248$; see Figure 4. The maximum likelihood parameter estimates before and after the estimated time of the change are $\hat{\theta}_b = 0.8515$ and $\hat{\theta}_a = 0.4087$, respectively. To obtain a confidence interval to the time of the change we have to simulate N samples of size 500 from the fitted bivariate Gumbel copula with a change in $k = 248$. Each sample is generated from

$$C^{Gu}_{\theta=0.8515}(\mathbf{x}_i)\mathbf{1}_{\{i\leq 248\}} + C^{Gu}_{\theta=0.4087}(\mathbf{x}_i)\mathbf{1}_{\{i>248\}}, \tag{14}$$

with $i = 1, \ldots, 500$. The frequency plot of the time of the changes estimated in the N simulations is in the left panel of Figure 6.

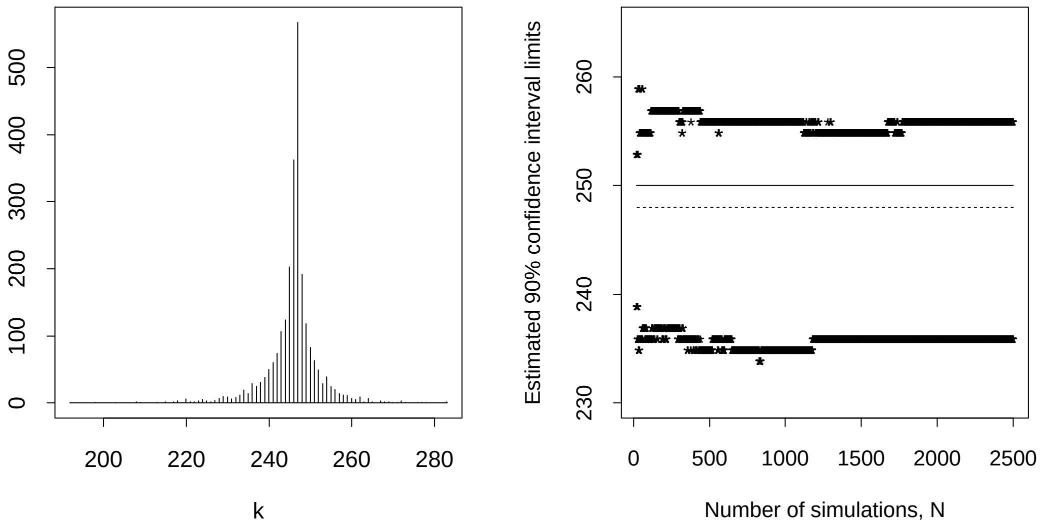

Figure 6: *Left panel: Frequency plot of the time of the change estimated from the simulated sets from the fitted distribution (14). Right panel:* 90% *bootstrap percentile confidence interval limits under resampling from the fitted model with one change for the data in the example from Section 2.2.*

In order to get a feeling for the number of bootstrap replicates needed to obtain a stable interval, in Figure 6 we plot the interval boundaries as a function of the number of replicates. In that plot the thin filled line is the change time in the data of the example ($k = 250$) and the dotted line is the estimated time of the change ($\hat{k}_{500} = 248$). We see that from around $N = 750$, the interval boundaries seem to stabilise around two values. We have that the 90% bootstrap percentile confidence interval for the time of the change given by (13) based on $N = 2'499$ resamples is

$$(\hat{k}^*_{500,((N+1)0.05)}, \hat{k}^*_{500,((N+1)0.095)}) = (236, 256).$$

2.5 Multiple Changes

The detection of several change–points in multidimensional processes with unknown parameters can be done using the so called binary segmentation procedure. This method was proposed by Vostrikova (1981) and enables to simultaneously detect the number and the location of the change–points. The method consists of first applying the likelihood ratio test from Section 2.1 for one change. If H_0 is rejected then we have the estimate of the time of the change $\hat{k}_n$. Next, we divide the sample in two subsamples $\{\mathbf{x}_i : 1 \leq i \leq \hat{k}_n\}$ and $\{\mathbf{x}_i : \hat{k}_n < i \leq n\}$ and test H_0 in each one of them. If we find a change in any of the sets we continue this segmentation procedure until we don't reject H_0 in any of the subsamples. In this paper, we will not discuss this procedure further.

3 A comment on pricing

From the previous sections we have seen that standard change–point procedures and bootstrap methodology can be combined to come up with estimates for change–point structures in copula data. Of course, combining the copula with certain marginal distributions, this work can be extended to more general bivariate (or indeed multivariate) models on $\mathbb{R}^2$ (respectively $\mathbb{R}^m$). One of the reasons that change–point analysis in dependence structures is important is because such changes often come from transitions from "normal" to "extreme" market conditions. Products priced correctly for instance for the former may be severely mispriced for the latter. In this section, and based on the example from Section 2.2, we will give a simple illustration of this. For further examples of applications of copula models as stress scenarios for financial positions in insurance or finance, see for instance Blum et al. (2002).

We will consider bivariate risks (U, V) following three possible bivariate distributions. First we take the copula model from Section 2.2 with the change–point, and consider the two resulting models before and after the change. Recall that the estimated Gumbel parameters were $\hat{\theta}_b = 0.8515$ before and $\hat{\theta}_a = 0.4087$ after the change–point. The homogeneous model ignoring the change–point yielded an estimate $\hat{\theta}_0 = 0.6325$. As typical payout functions we take

$$\Psi_1(U,V) = (U + V - 1.5)_+ \quad \text{and} \quad \Psi_2(U,V) = (U+V).\mathbf{1}_{\{U>0.8,V>0.8\}}.$$

We can estimate the expected value of these two positions, $E(\Psi_1)$ and $E(\Psi_2)$, using simulation. Table 4 has the Monte Carlo estimates for these expected values under the three models. Each value was obtained from a simulated set of $1'000'000$ pairs from the corresponding Gumbel copula. In brackets are the standard errors of the Monte Carlo estimations.

	Before time change $\hat{\theta}_b = 0.8515$	After time change $\hat{\theta}_a = 0.4087$	Ignoring change–point $\hat{\theta}_0 = 0.6325$
$\hat{E}(\Psi_1)$	0.0315 (0.0001)	0.0556 (0.0001)	0.0453 (0.0001)
$\hat{E}(\Psi_1)/\hat{E}_{\hat{\theta}_0}(\Psi_1)$	0.70	1.23	1
$\hat{E}(\Psi_2)$	0.1245 (0.0005)	0.2646 (0.0006)	0.1972 (0.0006)
$\hat{E}(\Psi_2)/\hat{E}_{\hat{\theta}_0}(\Psi_2)$	0.63	1.34	1

Table 4: *Monte Carlo analysis for two positions on the (U,V) simulated data, considering the change–point or ignoring that change.*

The estimated expected values vary considerably with the parameter of the distribution. From the table we can see that $\hat{E}(\Psi_1)$ after the time of the change is 1.23 of the value obtained if we ignore the change when performing the fitting. In the case of the second payout, the relative difference is even bigger, 1.34, between the estimated expected values considering and ignoring the change. This example illustrates the improvement that can be obtained when one recognises that there is a change–point in the data and one is able to take this change into account when fitting the data.

4 An example with insurance data

The Danish fire data is a set of $n = 2'493$ trivariate observations consisting in losses to buildings, losses to contents and losses to profits. These data were analysed in a one-dimensional setting for instance in Embrechts et al. (1997) and McNeil (1997). For bivariate analysis, see Blum et al. (2002). The independence between the n observations seems to be an acceptable assumption. In our example we use the variables losses to contents and losses to profits and only the observations with strictly positive components. The resulting set has $n = 517$ bivariate observations. An analysis of this data set reveals that the Gumbel copula yields a good fit; see Blum et al. (2002). The parameter estimate of the copula is $\hat{\theta} = 0.5385$ and a Kolmogorov-Smirnov goodness of fit test gives that we do not reject the model with a p–value of 0.52.

Here we are interested in testing for the existence of changes in the parameter of the dependence structure. If we test for one change using (1) we obtain that $Z_{517} = \max_{1 \leq k < 517}(-2\log(\Lambda_k)) = 3.4364$ which is attained at $k = 2$. The square root $Z_{517}^{1/2} = 1.8537$ lies out of any rejection region for the usual levels. Actually, we could only reject H_0 for levels $(1 - \alpha)$ smaller than 19.30%, using (3). In Figure 7 we plot the $(-2\log(\Lambda_k))$ for these data. The fact that the maximally selected likelihood ratio Z_{517} is attained at $k = 2$, as we saw in result (6), gives more confidence to not reject the no–change hypothesis. This analysis strengthens the approach taken in Blum et al. (2002) where functionals of the Danish data were priced using a homogeneous (i.e. no–change) Gumbel model.

5 Conclusion

Within finance and insurance, a considerable amount of effort is more recently put on the modelling of multivariate (typically dependent) data and the pricing of instruments based on such data. Copula based models are widely accepted as providing an interesting approach for the construction of such models looking at dependence beyond linear correlation. In this paper, we have looked at model changes within a parametric copula set up. Besides

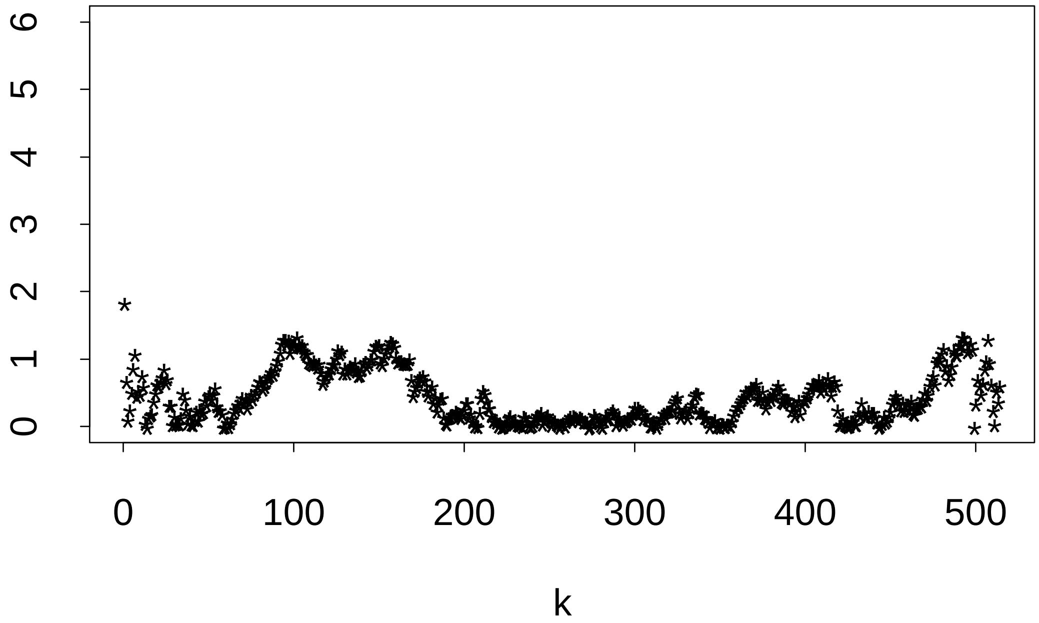

Figure 7: *Values of* $(-2\log(\Lambda_k))^{1/2}$*,* $1 \leq k < 517$*, for the Contents and Profits losses of the Danish fire data.*

explaining how the classical literature can be adapted to detect change–points in copula (hence dependence) parameters, we have also discussed the construction of bootstrap confidence intervals. We have illustrated the methodology introduced on simulated and real data. Several questions relevant for finance and insurance applications need further discussion. For instance, change–point analysis in dynamic finance data taking intertemporal dependence into account. Also, regression based models linking the change–point event to exogenous economic covariables can be worked out; this is especially important as behind every significant change–point in econometric data, are specific economic events acting as triggering events. We shall return to some of these questions in later publications.

References

Blum, P., Dias, A., and Embrechts, P. (2002). The ART of dependence modelling: the latest advances in correlation analysis. In: *Alternative Risk Strategies* (Morton Lane, Ed.) Risk Books, London, 339–356.

Boyer, B. H., Gibson, M. S., and Loretan, M. (1999). Pitfalls in tests for changes in correlations. *Board of Governors of the Federal Reserve System, International Finance Discussion Papers*, 597.

Csörgő, M. and Horváth, L. (1997). *Limit Theorems in Change–Point Analysis.* Wiley, Chichester.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and Their Applications.* Cambridge University Press, Cambridge.

Efron, B. and Tibshirani, R. J. (1993). *An Introdution to the Bootstrap.* Chapman & Hall, New York.

Embrechts, P., editor (2000). *Extremes and Integrated Risk Management.* RISK Waters Group, London.

Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance.* Springer, Berlin.

Embrechts, P., Lindskog, F., and McNeil, A. J. (2000). Modelling dependence with copulas and applications to risk management. To appear in *Heavy-Tailed Distributions in Finance* (2002)(S.T. Rachev, Ed.) North Holland, Amsterdam.

Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In: *Risk Management: Value at Risk and Beyond* (M. Dempster, Ed.) Cambridge University Press, Cambridge, 176–223.

Gombay, E. and Horváth, L. (1996). On the rate of approximations for maximum likelihood test in change–point models. *Journal of Multivariate Analysis*, 56(1):120–152.

Gombay, E. and Horváth, L. (1999). Change–points and bootstrap. *Environmetrics*, 10:725–736.

Hall, P. (1992). *The Bootstrap and Edgeworth Expansion.* Springer, New York.

Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.

Lehmann, E. L. (1991). *Theory of Point Estimation*. Wadsworth & Brooks, Pacific Grove, California.

Longin, F. and Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, LVI(2):649–676.

Loretan, M. and Phillips, P. C. B. (1994). Testing the covariance stacionarity of heavy–tailed time series: An overview of the theory with applications to several financial datasets. *Journal of Empirical Finance*, 2:211–248.

McNeil, A. J. (1997). Estimating the tails of loss severity distributions using extreme value theory. *Astin Bulletin*, 27(1):117–137.

Nelsen, R. B. (1999). *An Introduction to Copulas*. Springer, New York.

Vostrikova, L. J. (1981). Detecting "disorder" in multidimensional random processes. *Soviet Mathematics Doklady*, 24(1):55–59.