The Econometrics of Copula Functions

Mark Salmon

Mark.Salmon@wbs.ac.uk
Financial Econometrics Research Centre
Warwick Business School

http://www2.warwick.ac.uk/go/ferc/

Summer School: Frontiers in Financial Markets Mathematics
Copula Methods in Finance
University of Bologna
September 13-15th 2004

Financial Econometrics Research Centre WARWICK BUSINESS SCHOOL

Outline

- **Testing for Elliptic Distributions**
    - **Do we really need to worry about using Copulae at all?**
- **Estimating Parametric and Non-parametric Copulae**
    - **Maximum Likelihood;**
        - **Full MLE or sequential / concentrated ML**
    - **Semi-parametric estimation**
        - **Impact of the misspecification of marginal distributions on the estimation of copula parameters and measures of dependence**
    - **Non-parametric methods**
- **How to choose the Correct Copula?**
    - **AIC**
    - **Goodness of Fit Tests**
    - **Simulated Cox Tests**
- **Copula Quantile Regression and Tail area Dependence**
    - **Quantile conditional relationships implied by a given copula**
    - **Quantile dependency measures**

Testing for Elliptic Distributions

- A *spherical* distribution is an extension of the multivariate normal distribution $N_d(0, I)$ where $X = RU$ for some positive random variable R independent of a random vector U which is uniformly distributed on the unit hypersphere

$$\Omega_d = \{z \in \mathbb{R}^d \mid z^T z = 1\}$$

- An *elliptic* distribution is an extension of $N_n(\mu, \Sigma)$
- $N_d(\mu, \Sigma)$ is the distribution of $X = \mu + AY$

  where $Y \sim N_d(0, I)$ with the shape matrix $\Sigma = AA^T$

- So $X = \mu + RAU$ follows an elliptic distribution and the covariance matrix of X , Σ_0 is proportional to Σ
- An elliptic distribution is uniquely determined by its mean, covariance structure and its marginal distributions (ie. t_ν or normal,…)

- The copula of a multivariate elliptic distribution is uniquely determined by its correlation matrix and its type.

  eg. Gaussian Copula: $C_\rho^G(u) = \Phi_\rho(\phi^{-1}(u_1),\ldots,\phi^{-1}(u_n))$

  or

  Multivariate t_ν Copula: $C_{\nu,\rho}^{t}(u) = \Theta_{\nu,\rho}(t_\nu^{-1}(u_1),\ldots,t_\nu^{-1}(u_n))$
- So fat tails can be accommodated as a deviation from Gaussianity within the elliptic class although the use of the t Copula implies symmetry.
- The t Copula implies equal upper and lower tail dependence.

  $\lambda_u = 2\bar{t}_{\nu+1}\left(\sqrt{\nu+1}\sqrt{1-\rho}/\sqrt{1+\rho}\right)$ with

  $\bar{t}_{\nu+1}(x) = 1 - t_{\nu+1}(x)$

Manzotti et al. (2002), Jnl of Multivariate Analysis[1].

- Let $X_1,\dots X_n$ be an i.i.d. sample from a d dimensional distribution and we are interested in testing the null hypothesis that the sample comes from an elliptic distribution.
- Let $\bar{X}$ and S denote the sample mean and covariance matrix so $Y_k = S^{-1/2}(X_k - \bar{X})$ are called the scaled residuals $k = 1,\dots n$ and $W_k = Y_k / \|Y_k\|$ are their projections onto the unit sphere.
- If X is elliptically symmetric then W is approximately uniformly distributed on Ω_d and this can be verified in several ways—Manzotti et al. take the following route based on averaging spherical harmonics over the W_k's …..

[1] **see also Breymann, Dias and Embrechts, 2004, ETHZ, Zurich**

- Spherical harmonics in the 2 dimensional case are just trigonometric functions on the unit circle.
- Consider ε >0 fixed and let n_ε be the nearest integer to the left of εn. Let q_n be the ε empirical quantile for the radial variables $\|Y_1\|, \|Y_2\|, \ldots, \|Y_n\|$, n_ε are less than or equal to q_n and the rest are larger. We denote by $Q_n(h)$ the average over those W_k 's for which $\|Y_k\| > q_n$ of a function h defined on Ω_d: i.e.

$$Q_n(h) = \frac{1}{n} \sum_{k \le n} h(W_k) I_{\{\|Y_k\| > q_n\}}$$

and the statistic is then given by

$$Z_n^2 = \sum_{h \in \Im} Q_n^2(h)$$

where the function *h* selects the spherical harmonics of appropriate degree . eg. For the *d=2* case

$h_{1j}(W_k) = 2^{1/2}\cos(j\vartheta_k), h_{2j}(W_k) = 2^{1/2}\sin(j\vartheta_k)$ for $3 \le j \le 6$ and $W_k = (\cos(\vartheta_k), \sin(\vartheta_k))$.

- Manzotti et al showed that the asymptotic distribution of Z_n^2 is then chi squared regardless of the parameters defining the underlying distribution.
- Statistics related to Z_n^2 have shown very good power in testing for multivariate normality against a wide range of alternatives (See Manzotti and Quiroz (2001),TEST)

An Example

- We will look at the question of the information content of realised volatility and volume in terms of explaining daily returns.
- The Mixture of Distributions hypothesis (Clark 1972 and many papers since) suggests the non-normality of returns follows from the irregular arrival of information events in the market- so the price change process is subordinated to the information arrival process which is random hence the central limit theorem does not apply and returns are Non- Gaussian.
- This subordination process effectively defines a new time scale- instead of clock time we define a new clock which advances time on the basis of market activity- as the market is more active time moves faster –when the market is slow time moves forward slowly.

- This leads naturally to Stochastic Volatility models- but questions still remain as to how to measure information flow into the market—is it the number of transactions in any interval, the volume or some other measure of activity.
- The importance of this issue is that if we can find the appropriate time deformation variable which measures market activity then by re-expressing returns in market time in place of clock time we can recover conditional-Gaussianity in returns.
- If we can do this satisfactorily across the assets we are interested in then standard finance theory and methods will apply and this in fact will lessen the demand for non-Gaussian methods such as Copula!
- We will consider the relationship between volume and realized volatility and returns.
- Realised volatility is a measure of volatility that has attracted a lot of interest in the literature recently—it is

essentially a daily measure of volatility made up from the sum of intraday volatilities.

- We use tick data on Dell aggregated up to the 5 minute level which we use to construct daily realised volatility- which we compare with daily volume of transactions and the daily returns.
- The figures below show some of the characteristics of the data.

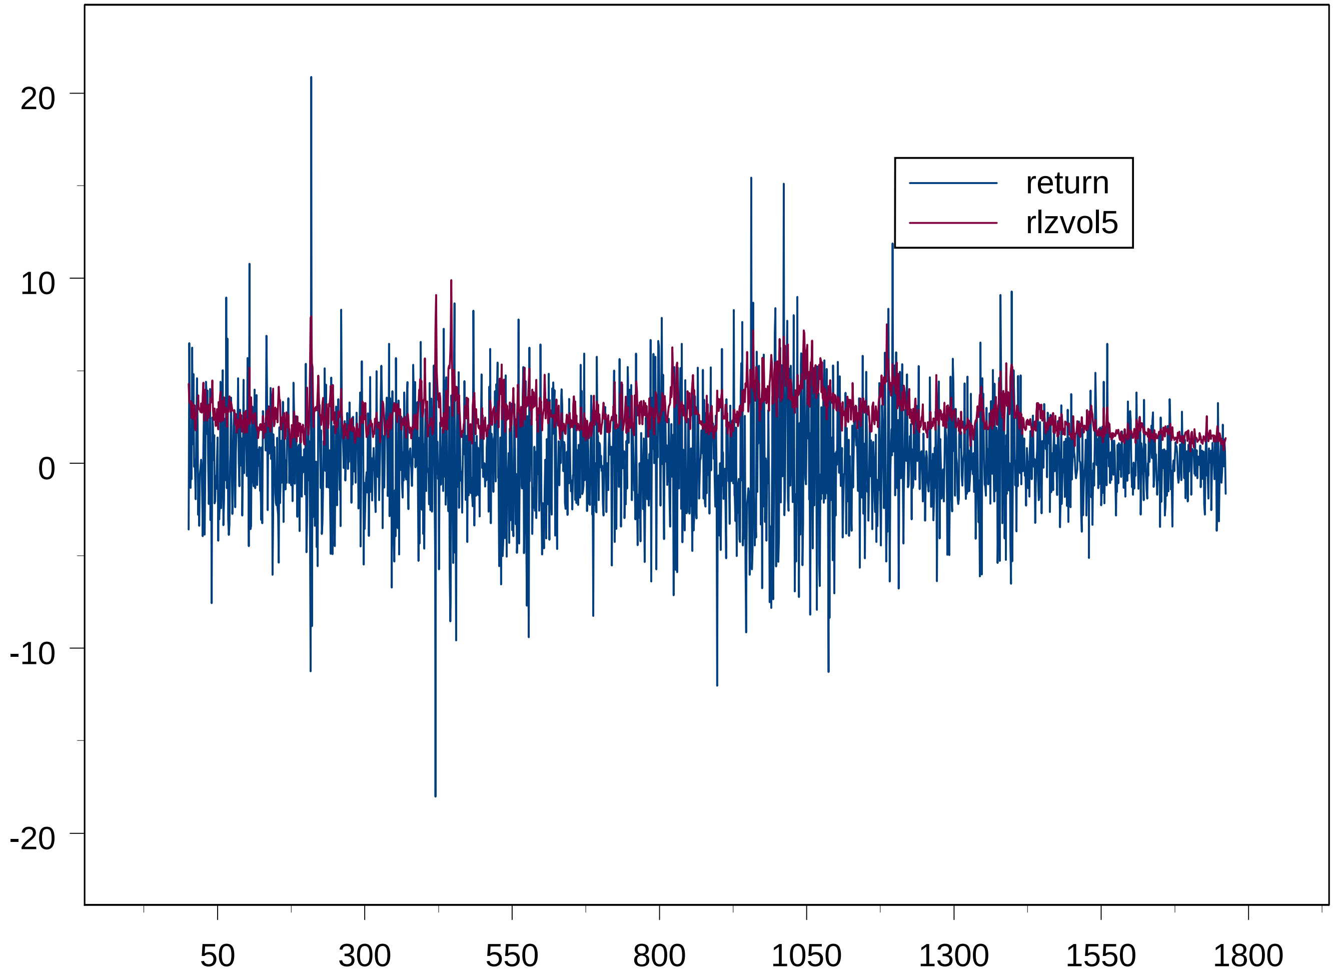

return
rlzvol5
20
10
0
-10
-20
50
300
550
800
1050
1300
1550
1800

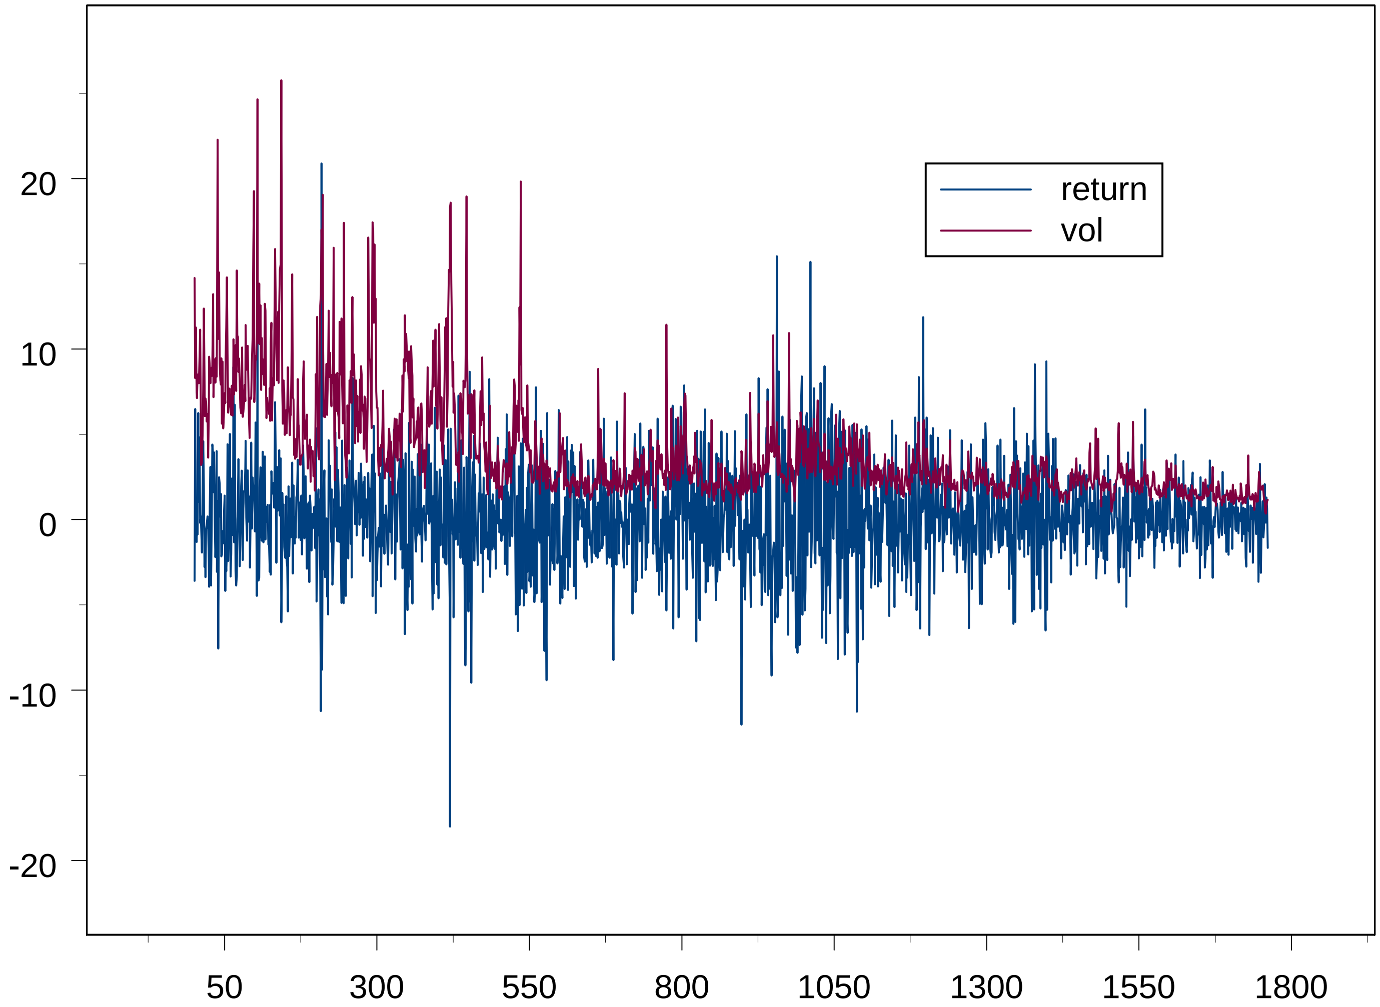
return
vol
20
10
0
-10
-20
50
300
550
800
1050
1300
1550
1800

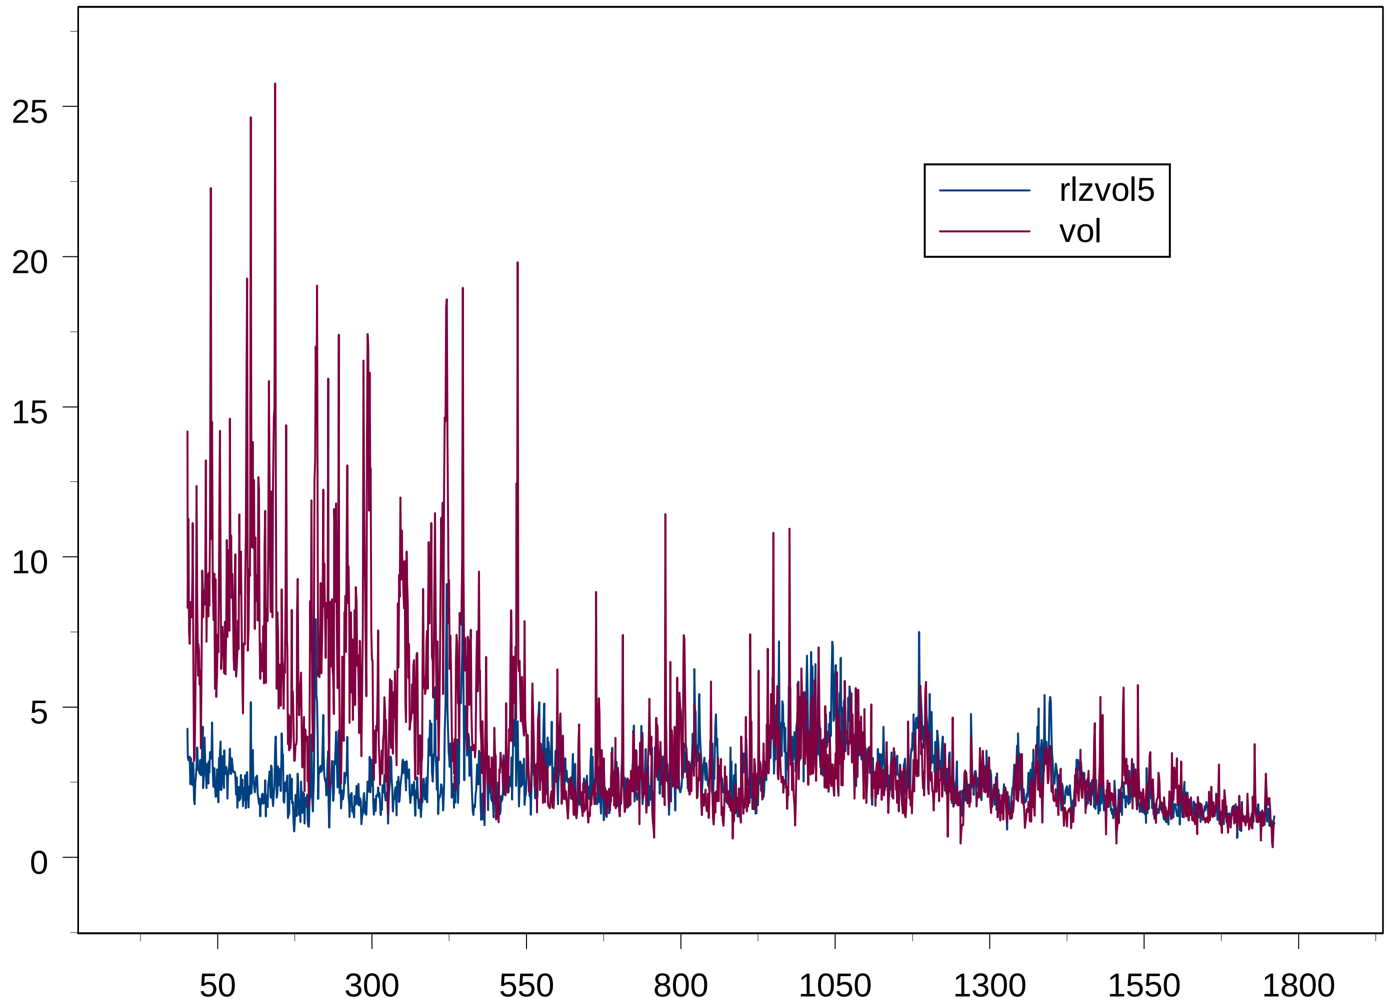
rlzvol5
vol
25
20
15
10
5
0
50
300
550
800
1050
1300
1550
1800

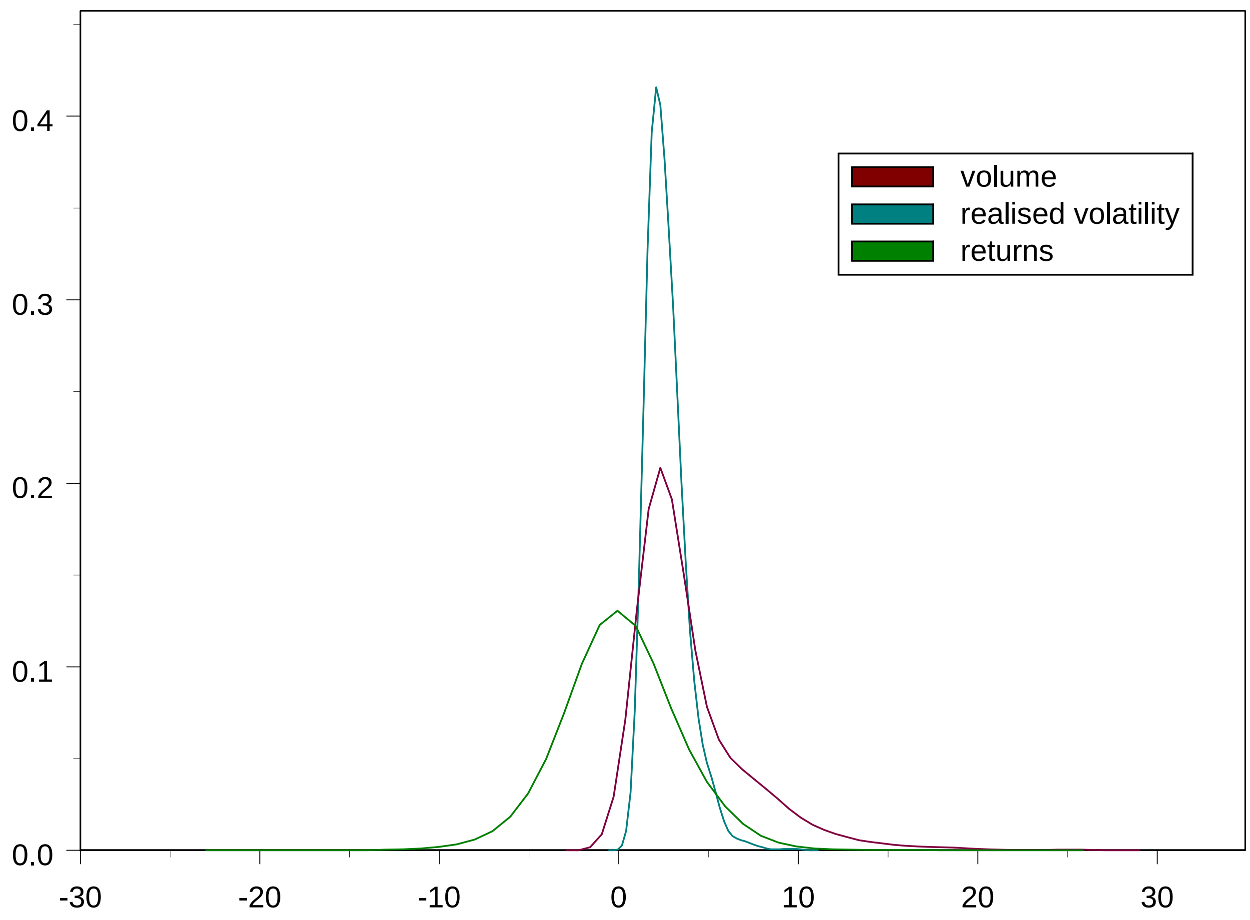
volume
realised volatility
returns
0.4
0.3
0.2
0.1
0.0
-30
-20
-10
0
10
20
30

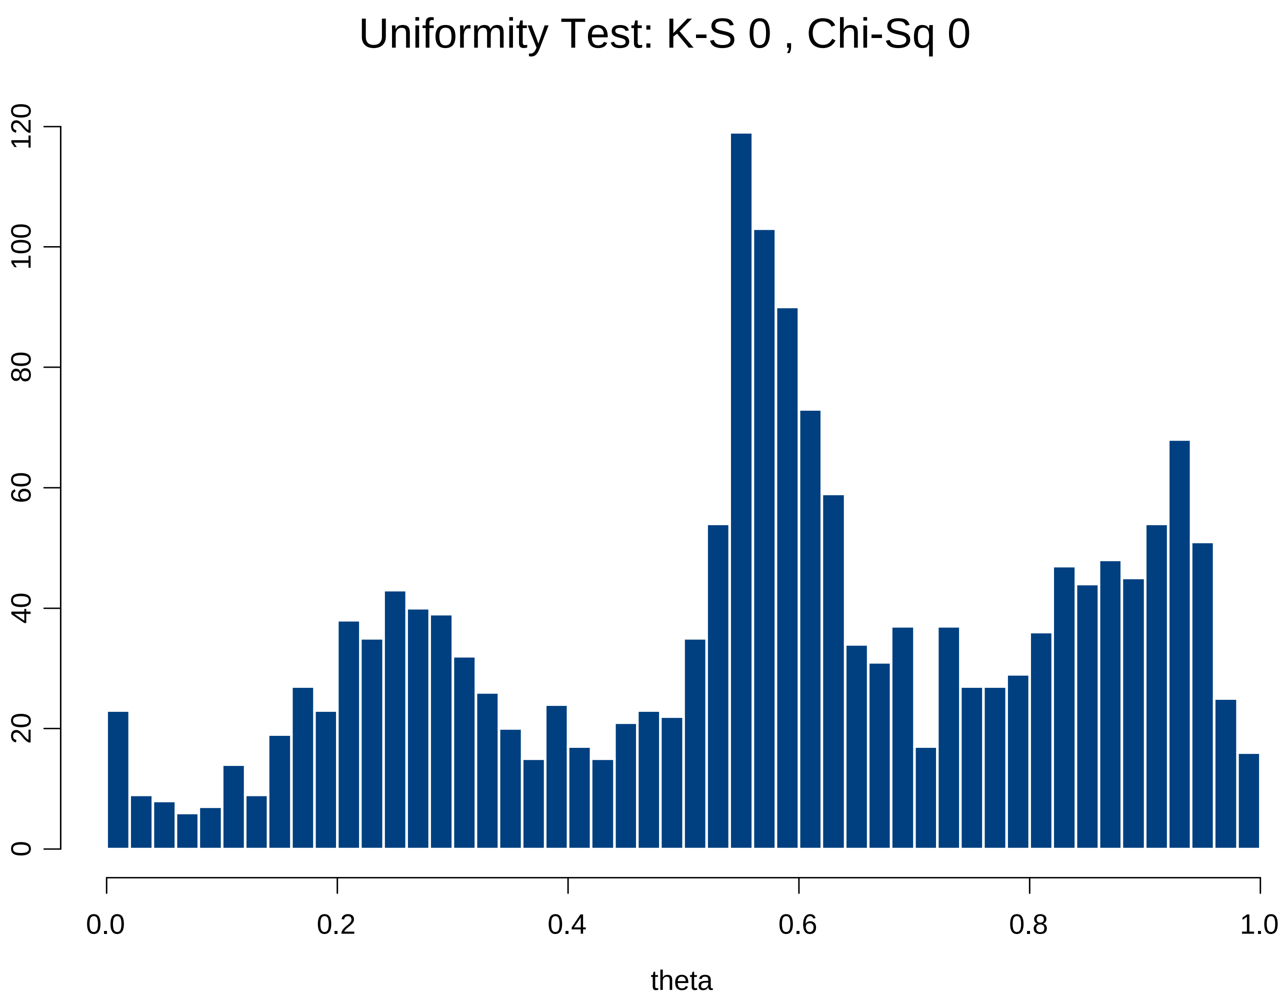

Volume: Realised Volatility

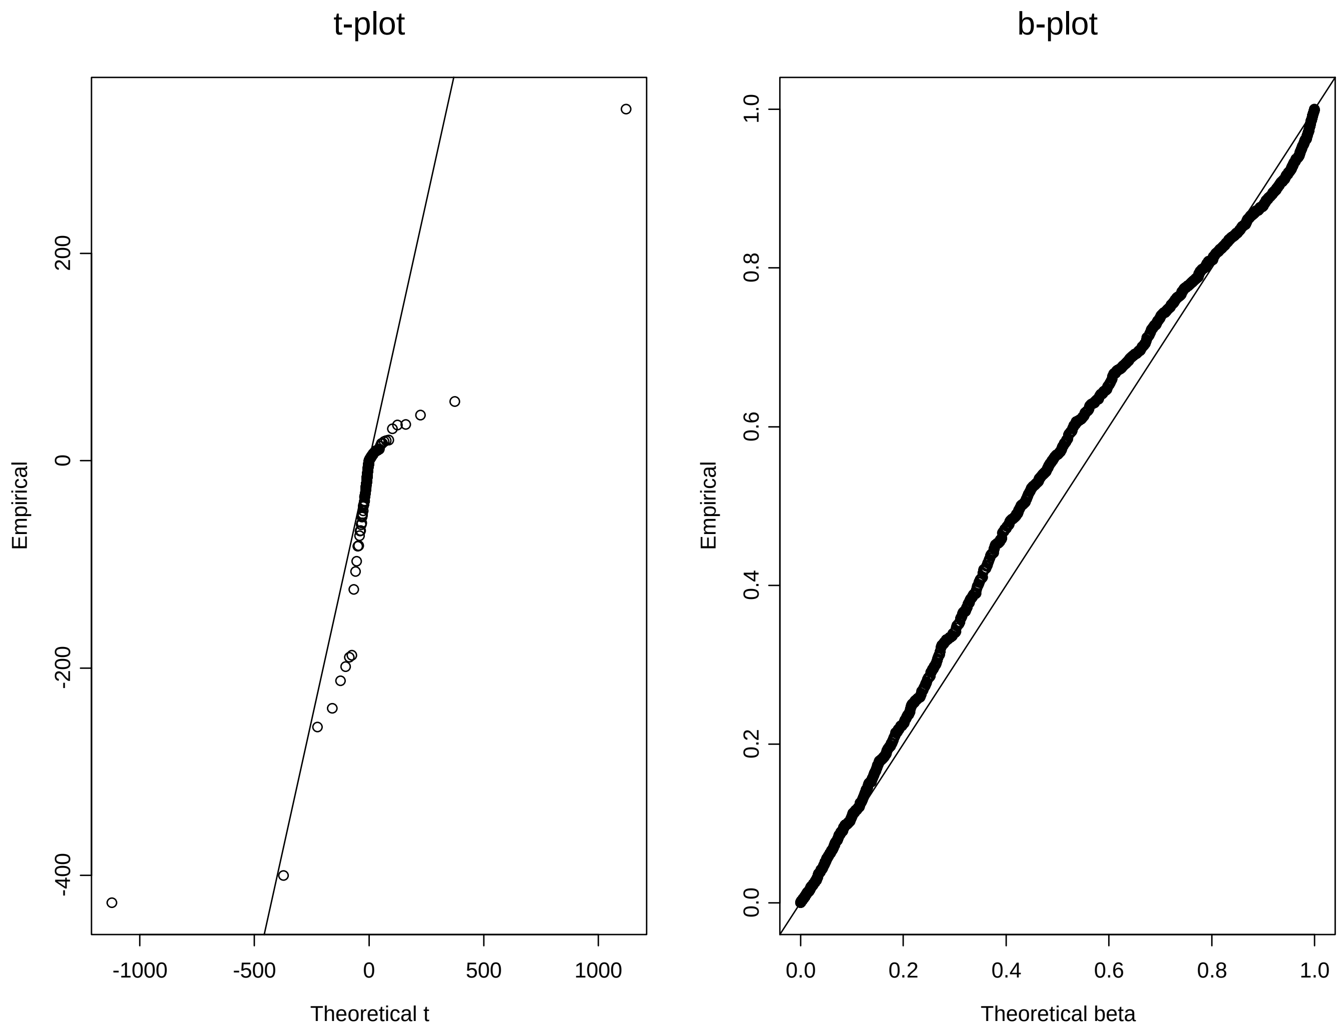

QQ Plot for Ellipticity Volume+Realised Volatility

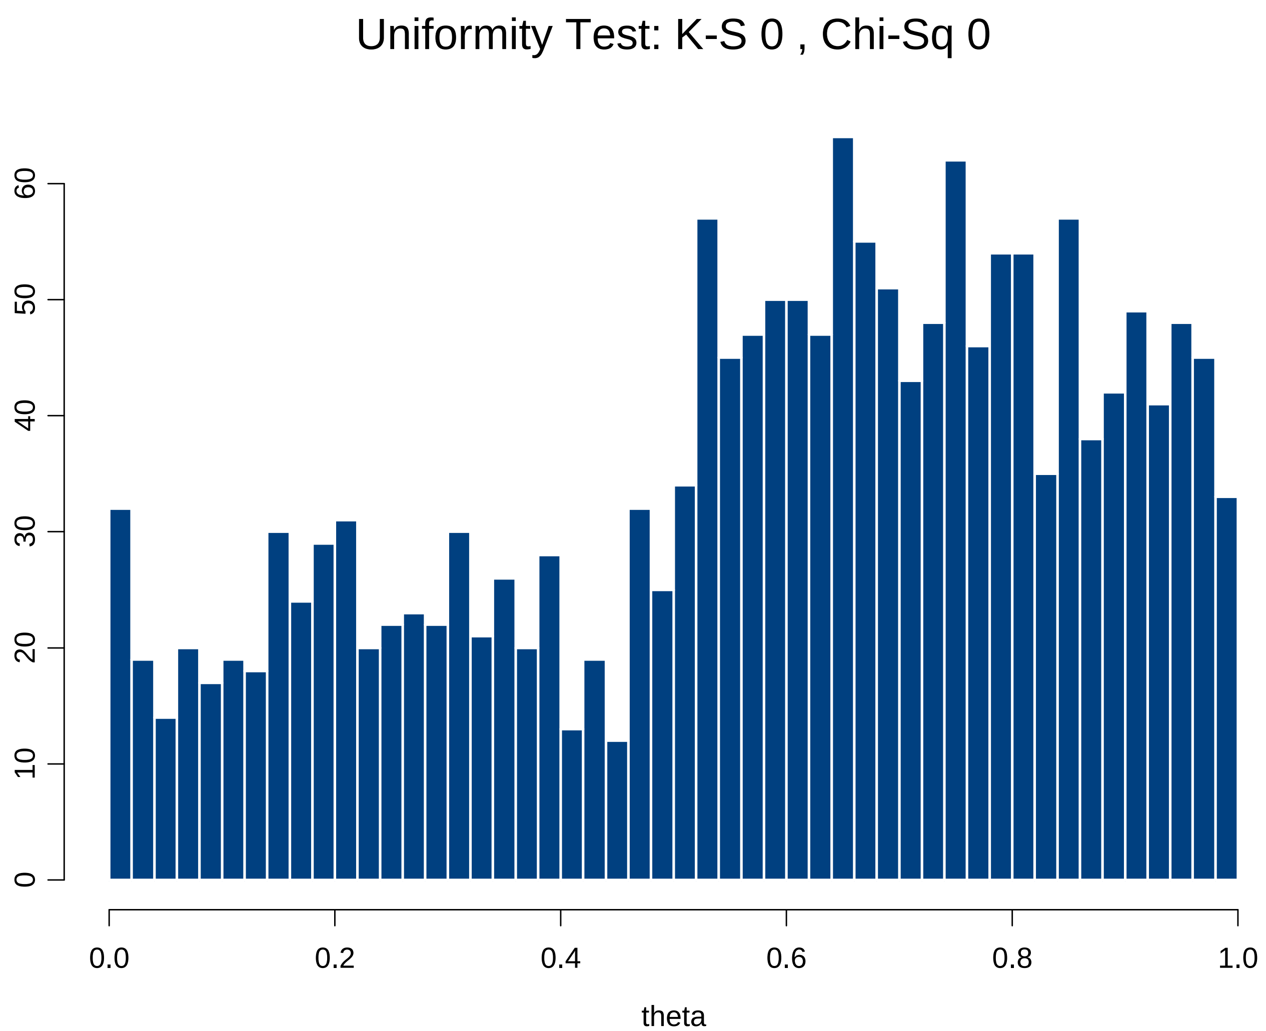

Returns- Volume

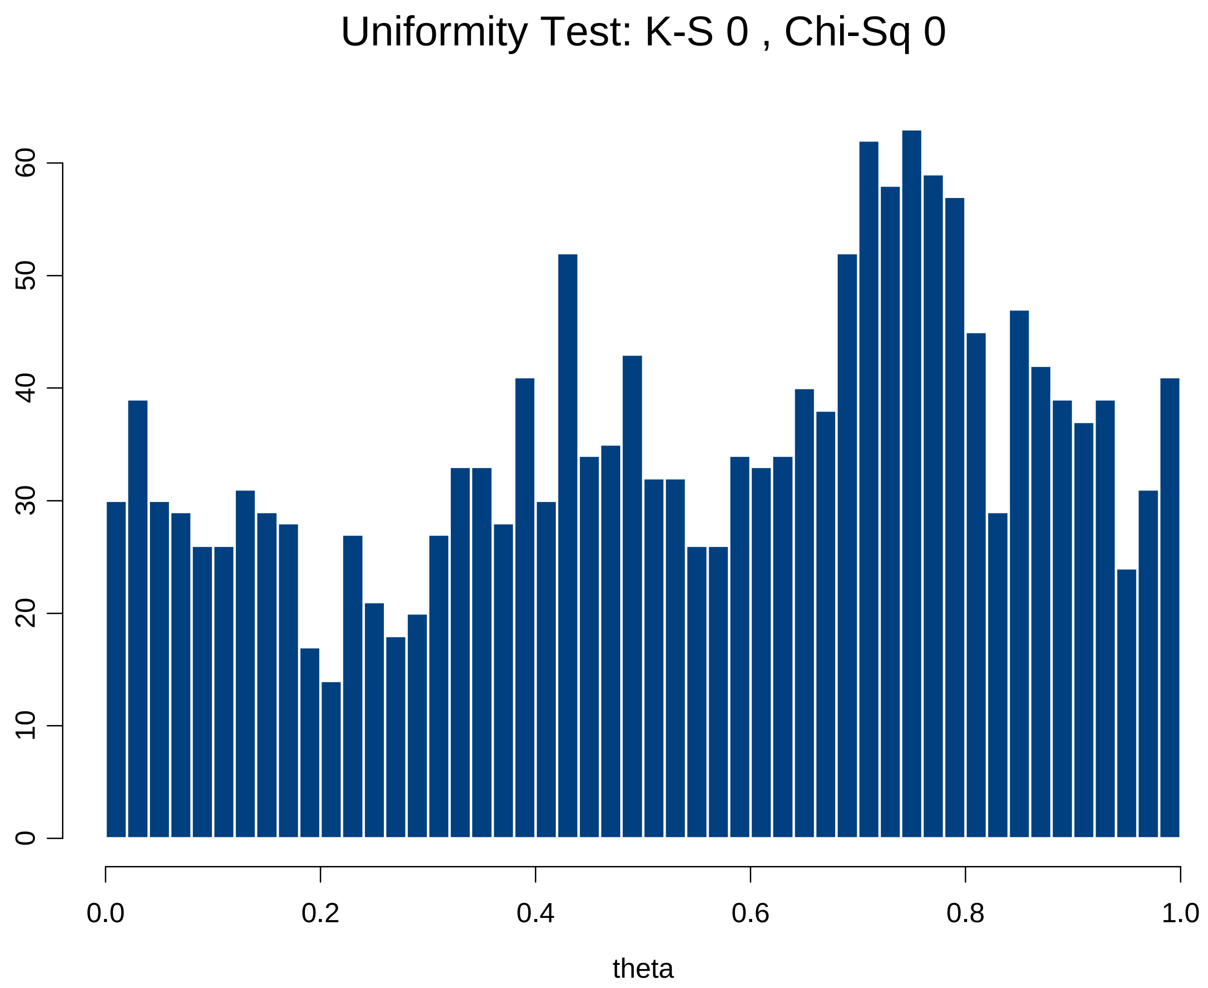

Returns- Realised Volatiltiy

Estimating Parametric and Non-parametric Copulae

- We assume a d dimensional cumulative distribution expressed in its Copula form as

$$F(x) = C_{\vartheta}(F_1(x_1, \delta_1), \ldots F_d(x_d, \delta_d), \vartheta)$$

where C_{ϑ} is the copula function and the F_i are the margins.

- One of the main advantages of the Copula approach is that it enables us to construct the multivariate distribution by first modelling each univariate marginal distribution, either parametrically or non-parametrically and then, in a second step, specifying and estimating the Copula function which captures all the dependencies between the margins.
- This second step is in fact much more complicated than most people assume and has been relatively under researched.

- Part of the problem lies in dynamic dependencies, serial correlation or autoregressive heteroskedasticity in financial returns. So how can we can we estimate Copulae?
- As long as the true copula belongs to a parametric family

$$\mathbf{C} = \left\{ C_{\vartheta}, \vartheta \in \Theta \right\}$$

then standard application of Maximum likelihood methods will, in principle, deliver consistent, efficient and asymptotically normally distributed estimates.

- Two routes have been followed;

a) Parametric

1. Concentrated MLE (or IFM)
2. Full MLE

b) Semi-parametric

3. Canonical MLE (CML)

a) Parametric MLE:

The density function of X is written

$$f(x;\delta_1,\delta_2,....,\delta_d,\vartheta) = c(F_1(x_1,\delta_1),...F_d(x_d,\delta_d);\vartheta)\prod_{i=1}^{d} f_i(x_i,\delta_i)$$

where the copula density is given by

$$c(u_1,u_2,..,u_d,\vartheta) = \frac{\partial^d C(u_1,u_2,..,u_d,\vartheta)}{\partial u_1 \partial u_2 ... \partial u_d,} \quad .$$

The log-likelihood is then given by

$$l(\vartheta,\delta;X) = \sum_{t=1}^{n}\left(\log c(F_1(x_{1t},\delta_1),...F_d(x_{dt},\delta_d);\vartheta) + \sum_{i=1}^{d} \log f_i(x_{it},\delta_i) \right)$$

Note that the estimation of the copula relies on the specification of parametric univariate marginal distributions in some way – once this has been done - their parameters δ_i can either be

1. estimated prior to inclusion in the full likelihood – a two step procedure-(Concentrated MLE) and substituted into the full likelihood function; so first estimate the δ's by numerically maximising $\sum_{t=1}^{n} \log f_i(x_{i,t}, \delta_i) \quad \text{for } i = 1..d$

and then maximising

$$l(\vartheta; X) = \sum_{t=1}^{n} \left(\log c(F_1(x_{1t}, \hat{\delta}_1), \ldots F_d(x_{dt}, \hat{\delta}_d); \vartheta) \right)$$

to get the estimates of the copula parameters ϑ.

2. alternatively the marginal parameters δ can be estimated at the same time as the copula parameters ϑ are estimated, although as the scale of the problem increases this option is often considered too burdensome computationally. Either way the prior specification of the marginal distributions leads to potential misspecification problems

- The properties of the Full Maximum Likelihood estimators under 2. above are well established under a range of regularity conditions and the procedure delivers consistent, efficient and asymptotically normal estimators when the model is correctly specified.

$$\sqrt{n}(\hat{\vartheta}-\vartheta)\to_d N(0,\Im^{-1}(\vartheta))$$

- Estimation by the two step procedure generally delivers less efficient but consistent estimates of the copula parameters and the finite sample estimates and properties will be different as different likelihood equations are used to construct the estimates. See Newey Ec.Letters 1984 for a Method of Moments interpretation of sequential estimators.

- When there is misspecification in some form then we should use, on a regular basis, the theory outlined by Hal White (Econometrica 1982) for Maximum Likelihood Estimation in Misspecified models should be adopted. This implies that the ML estimates will converge to the values that are closest to the true values within the assumed model space in the Kullback Leibler sense of distance. Fisher's information matrix is no longer the correct basis for constructing confidence intervals. This issue appears to have been almost completely ignored in the copula literature.
- This is particularly important in this area as it has been common practice to simply assume a copula form or to assume the form of convenient margins.

b) Semi-Parametric:

The semi-parametric route avoids the problem of specifying the parametric form for the margins and instead substitutes the univariate empirical cumulative distribution functions for the margins. So for each *i* we substitute

$$\hat{F}_i(x) = \frac{1}{n+1}\sum_{j=1}^{n} I_{\{X_{ij}<x\}}$$

and maximise the "pseudo log likelihood"

$$\sum_{j=1}^{n} \log c_{\vartheta}(\hat{F}_1(x_{1,j}),\ldots\hat{F}_d(x_{d,j}))$$

However this approach leads to inefficient parameter estimates *except* in the case of the normal copula and independence.

- Essentially the marginal distributions appear statistically as infinitely dimensioned nuisance parameters.

- In both cases the copula function itself maybe misspecified and hence as in standard likelihood methods the asymptotic variance should be replaced by the sandwich estimator as in White (Econometrica,1982) or Gourieroux et al Pseudo MLE(Econometrica,1984)
- Fermanian and Scaillet (2004) provide a Monte Carlo example of the impact of misspecification of margins.
- The true model corresponds to a Frank copula with two student t margins, the copula parameter θ=1 and 2 with degrees of freedom equal to 3 for both margins. The misspecified model assumes Gaussian margins. The results for four estimators are reported, full MLE, the concentrated 2

step MLE, the semi- parametric method and MLE applied to the true model. Nreps =1000.

	Sample Size : n=200			
θ=1	**Full MLE**	**Two step**	**Semi-par**	**true**
Bias	0.7012	0.6094	-0.0206	-0.0164
MSE	1.5119	1.0591	0.1754	0.1797
θ=2				
Bias	1.1144	0.9292	-0.0202	-0.0124
MSE	2.2869	1.4851	0.1913	0.1928

	Sample Size : n=500			
$\theta=1$	**Full MLE**	**Two step**	**Semi-par**	**true**
Bias	0.7720	0.6931	-0.0081	-0.0067
MSE	1.1114	0.8184	0.0673	0.0684
$\theta=2$				
Bias	1.2165	1.0354	-0.0067	-0.0045
MSE	2.0494	1.3977	0.0730	0.0729

	Sample Size : n=1000			
$\theta=1$	**Full MLE**	**Two step**	**Semi-par**	**true**
Bias	0.7904	0.7190	-0.0046	-0.0048
MSE	0.9647	0.7397	0.0360	0.0360
$\theta=2$				
Bias	1.2553	1.0784	-0.0037	-0.0037
MSE	1.9702	1.3853	0.0387	0.0382

- Clearly the effects of misspecifying the marginal distributions can be a severe bias and this in turn leads to an over-estimation of the degree of dependence in the data.
- While asymptotically they have the same distribution in this the bias seems to be consistently larger in the one step method.
- The MSE is also considerably larger when compared to the semi-parametric method (which by assumption is in fact correctly specified)
- These same results have also been found by Silvapulle,Kim,Silvapulle, Monash WP-2004.

It would seem natural to follow the semi-parametric route…but

- As mentioned above; following the original suggestion of the method by Oaks (1994), estimation by the Semi-parametric method was considered by Genest,Ghoudi and Rivest (1995, Biometrika) who show only consistency and asymptotic normality of the resulting estimates.

- Genest and Werker (2002, proceedings of Conference on Distributions with Given Marginals and Statistical Modelling ed C.M.Cuardas and J.A. Rodrigues- Lallena) showed that the standard estimator is in general semi-parametrically inefficient. Using standard empirical distribution functions for the margins leads to inefficient estimates and semi-parametrically efficient estimates of the margins need to be used.

- I know that Yanqin Fan and Xiaohong Chen at Vanderbilt and John Einmahl and Bas Werker at Tilburg are working on this issue but this is as far as it has got I believe.

- Yanqin Fan and X. Chen (NYU) (2004) have extended the semi-parametric estimator to time series models (non i.i.d data and shown the consistency and asymptotic normality under β mixing processes. They also provide a consistent estimate of the asymptotic variance using a parametric bootstrap procedure.

- Bouye, Gaussel and Salmon (2003) FERC consider semiparametric nonlinear dynamic models implied by different copula and develop auto-concordance measures in place of autocorrelation.

Mispecification Chen Fan paper

Fully Non-Parametric Estimation:

- Since the Copula is just a multivariate distribution function standard smoothing methods can be used to estimate Copula non-parametrically using Kernel density estimators.
- Deheuvals (1978)(1981) proposed an early nonparametric estimator which has become know as the empirical copula but these are highly discontinuous and not of much use in practice.
- Fermanian and Scaillet (2003) WP HEC Geneva propose nonparametric estimators for copulas of time series for multivariate stationary processes satisfying strong mixing conditions.

- The pdf of X_{jt} at x_j i.e. $f_j(x_j)$ is estimated by the usual kernel estimator

$$\hat{f}_j(x_j) = (nh_j)^{-1} \sum_{j=1}^{n} k_j \left(\frac{x_j - X_{jt}}{h_j} \right)$$

with the usual kernel function conditions such that

$$\int k_j(x)dx = 1 \quad j = 1, \dots n$$

while the pdf of $\mathbf{X_t}$ at $\mathbf{x}$, $f(\mathbf{x})$ will be estimated by

$$\hat{f}(\mathbf{x}) = (T|h|)^{-1} \sum_{j=1}^{n} k(\mathbf{x} - \mathbf{X}_t; h)$$

hence an estimator of the cumulative distribution of $\mathbf{X}_t$ at some point $\mathbf{x}$ is given by

$$\hat{F}(\mathbf{x}) = \int_{-\infty}^{x_1} \dots \int_{-\infty}^{x_d} \hat{f}(\mathbf{x}) dx$$

and if a single Gaussian Kernel is used k_j (x)=φ(x) then we have

$$\hat{F}(\mathbf{x}) = \mathbf{n}^{-1} \sum_{j=1}^{n} \prod_{i=1}^{d} \Phi((x_i - X_{it})/h_i)$$

so the estimate of the copula at some point is simply given by

$$\hat{C}(\mathbf{u}) = \hat{F}(\hat{\zeta}) \quad \text{where } \hat{\zeta} = \inf_{x \in \mathbb{R}} \left\{ x : \hat{F}_j \geq u_j \right\}$$

- Fermanian and Scaillet demonstrate the asymptotic normality of these kernel estimators for copulas under suitable regularity conditions and the asymptotic behaviour of the bandwidth and the mixing assumptions on the data.
- They then apply these estimators in a Monte-Carlo study for a VAR(1) with Gaussian innovations showing satisfactory bias and MSE and then consider empirical measures of

positive quadrant dependence and VaR for CAC40-DAX35 and SP500-DJ

How to choose the Correct Copula?

Akaike's Information Criterion:

- This is the simplest method that has been regularly used to select a copula up to this point. It does not provide a test and so we have no understanding of the size of the decision rule that is implied nor its power—it is simply a selection criterion that says we should use the copula that delivers the smallest value of the adjusted likelihood criterion;

$$AIC = -2L(\hat{\vartheta}, \hat{\delta}; x) + 2q$$

where q is the number of parameters of the family being fitted. The adjustment to the maximised likelihood penalizes models that employ more parameters.

Goodness of Fit tests:

- The probability integral transform (PIT) of Rosenblatt (1952) and can be used for any copula family. The basic idea follows from thinking about the transformation from the data to the cdf:

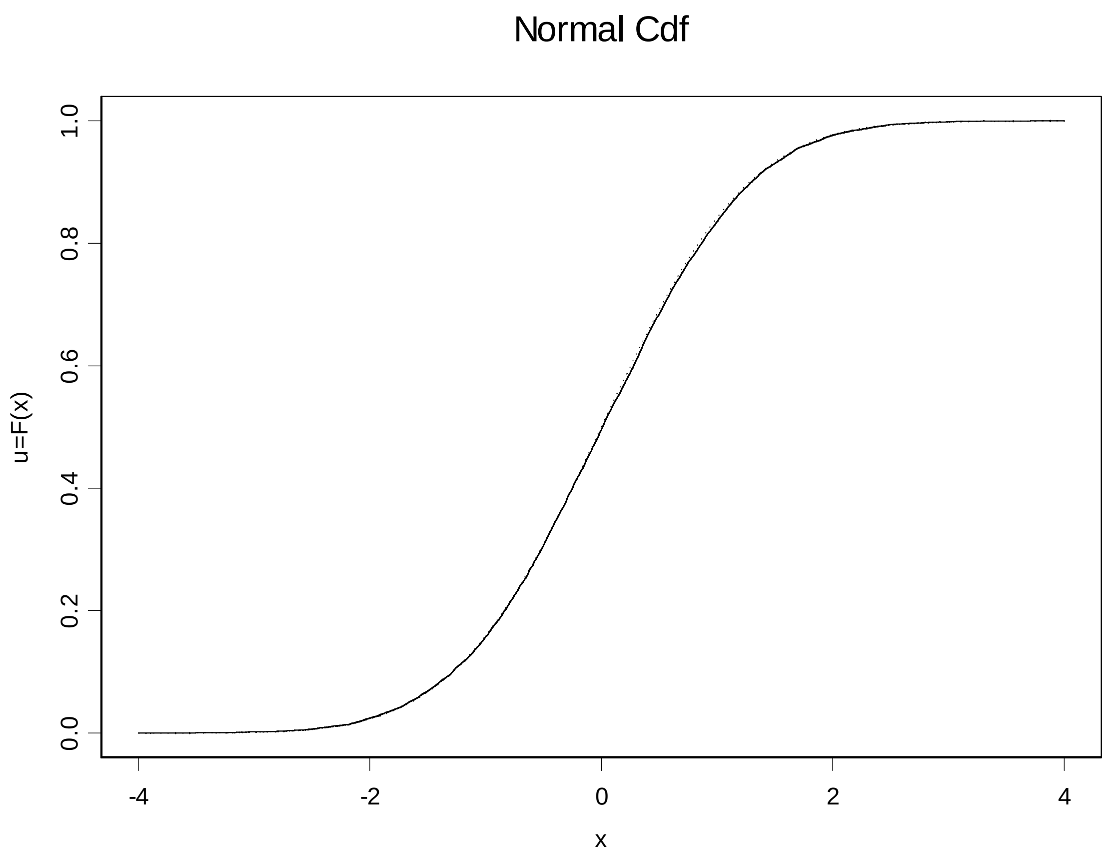

- Essentially if we define new variables, *z*, by transforming our observed data by using an assumed distribution –say *F(x)*- so we have $z = F(X)$ then if the data is actually drawn from the *F(.)* distribution so that $x = F^{-1}(u)$ then the *z* variables will be uniform and i.i.d..as $z = F(F^{-1}(u)) = u \ldots$
- If the data were generated by some other distribution, say *G(.)*, then this would not hold as we would have $z = F(G^{-1}(u)) \neq u$. This simple observation underlies most of the recent work in testing distributions and copula.
- Let $\{C_0(u_1,\ldots u_d;\vartheta\}$ be a class of parametric copula; we are interested in testing:

$$H_0 : \Pr(C(U_1,\ldots U_d) = C_0(U_1,\ldots U_d;\vartheta_0)) = 1 \text{ for some } \vartheta_0 \in \Theta$$

against

$$H_a : \Pr(C(U_1,\ldots U_d) = C_0(U_1,\ldots U_d;\vartheta)) < 1 \text{ for all } \vartheta \in \Theta$$

- In the univariate case a number of tests have been put forward which compare the empirical cdf with the cdf of the null distribution. In particular Kolmogorov-Smirnov and the Anderson-Darling Statistics;

1. Kolmogorov-Smirnov: $d_1 = \max_z \left| F_{z^2}(z^2) - F_{\chi^2}(z^2) \right|$

2. Average K-S: $d_2 = \int \left| F_{z^2}(z^2) - F_{\chi^2}(z^2) \right| dF_{\chi^2}(z^2)$

3. Anderson Darling: $d_3 = \max_z \dfrac{\left| F_{z^2}(z^2) - F_{\chi^2}(z^2) \right|}{\sqrt{F_{z^2}(z^2)\left[1 - F_{\chi^2}(z^2)\right]}}$

4. Average A-D: $d_4 = \int \dfrac{\left| F_{z^2}(z^2) - F_{\chi^2}(z^2) \right|}{\sqrt{F_{z^2}(z^2)\left[1 - F_{\chi^2}(z^2)\right]}} dF_{\chi^2}(z^2)$

Breymann, Dias, Embrechts(2004) use the PIT idea as follows:

5. Let $\mathbf{X} = (X_1, X_2, \ldots X_d)$ denote a random vector with cdf. $F_X(x_1, x_2, \ldots x_d)$ and let $F_{X_i}(x_i) = P(X_i \le x_i)$ be the distribution function for the univariate margins X_i , i=1,…,n.

6. Consider the following *PIT* defining *d* random variables $Z_i = T(X_i)$ which will be uniform and independently distributed:

$$T(x_1) = P(X_1 \le x_1) = F_{X_1}(x_1)$$

$$T(x_2) = P(X_2 \le x_2 \mid X_1 = x_1) = F_{X_2 \mid X_1}(x_2 \mid x_1)$$

$$\vdots$$

$$T(x_d) = P(X_d \le x_d \mid X_1 = x_1, \ldots, X_{d-1} = x_{d-1})$$
$$= F_{X_d \mid X_1 \ldots. X_{d-1}}(x_d \mid x_1, \ldots, x_{d-1})$$

Suppose that C is a copula so that

$$F_X(x_1, x_2, ..., x_d) = C(F_{X_1}(x_1), \ldots F_{X_d}(x_d))$$

and if $C_i(u_1, ... u_i)$ denotes the joint i-distribution of $(U_1, ..., U_i)$ with $C_1(u_1) = u_1$ and $C_d(u_1, ..., u_d) = C(u_1, ..., u_d)$ then the conditional distribution of U_l given the values of $U_1, ..., U_{i-1}$ is given by

$$C_i(u_i \mid u_1, ..., u_{i-1}) = \frac{\partial^{i-1} C_i(u_1, ..., u_i)}{\partial u_1 .... \partial u_{i-1}} \Big/ \frac{\partial^{i-1} C_{i-1}(u_1, ..., u_{i-1})}{\partial u_1 .... \partial u_{i-1}}$$

for i=2,….d.

7. Hence we can write the variables Z_i for i=2,….d. using the conditional distributions C_i as

$$Z_i = C_i(F_{X_i}(X_i) \mid F_{X_1}(X_1), ..., F_{X_{i-1}}(X_{i-1}))$$

8. As the copula is a multivariate distribution function it follows that H_0 holds if and only if the PIT transformed variables are i.i.d. and uniformly distributed $\left[0,1\right]$ hence H_0 is equivalent to

$$H_0^{'} : P(g(Z_1,...Z_d)-1)=1$$

where $g(z_1,...z_d)$ is the joint density of the transformed variables.

- Breymann et al.then use this construction as follows:-

If $(F_{X_1}(X_1), F_{X_2}(X_2),..., F_{X_d}(X_d))$ has distribution function C_0 then $\phi^{-1}(Z_i), i=1,...,d$ are i.i.d. $N(0,1)$.

So $S = \sum_{i=1}^{d} \left(\phi^{-1} \left(Z_i \right) \right)^2$ will have a chi squared distribution with d degrees of freedom.

In the case $d=2$ we have;

$$S(X_1, X_2) = (\phi^{-1}(F_{X_1}(X_1))^2 + (\phi^{-1}(C_2(F_{X_2}(X_2) \mid F_{X_1}(X_1))))^2$$

They then use the Anderson Darling Statistic to carry out the test.

- However the marginal distributions in their set up are unspecified, and should be treated as infinite dimensional nuisance parameters. As a results the existing critical values for the A-D are invalid and there is considerable uncertainty about the size of their test.

- Malvergne and Sornette suggest the use of Bootstrap critical values for the A-D test in this case.

Chen,Fan, Patton (2004)

- Propose two simple asymptotically valid tests for the goodness of fit for any parametric copula model of the contemporaneous dependence between two time series based on the multivariate PIT and kernel smoothing—both are asymptotically normally distributed under the null hence distribution free and easy to compute.
- The first is consistent but requires the kernel estimation of a multivariate density function and hence is suitable when only a small number of assets is being considered.
- The second may not be consistent in all directions away from the null- ie. all alternatives --but it only requires the kernel estimation of a univariate density and so is useful when testing the dependence between a large number of assets.

- Their approach is to test $H_0^{'}$ by first estimating the joint density $g(z_1,...z_d)$ by a multivariate kernel estimator

$$\overline{g}(z_1,...z_d) = \frac{1}{nh^d}\sum_{t=1}^{n}\left[\prod_{j=1}^{d} K_h(z_j, Z_{jt})\right]$$

and the test is based on $\overline{I}_n = \int_0^1 ....\int_0^1 \left[\overline{g}(z_1,...z_d) - 1\right]^2 dz_1...dz_d$ once the unobserved Z variables have been estimated from

$\hat{Z}_{it} = \hat{F}_1(X_{1t})$ and $\hat{Z}_{jt} = C_0(\hat{F}_j(X_{jt}); \hat{\vartheta} \mid \hat{F}_1(X_{1t})...\hat{F}_{j-1}(X_{j-1t}))$ j=2,…d

where $\hat{\vartheta}$ is a consistent estimator of ϑ under H_0 and

$$\hat{F}_j(x_j) = \frac{1}{n+1}\sum_{t=1}^{n} 1_{\{X_{jt} \le x_j\}}$$

Under suitable normalisation this statistic is shown to be asymptotically Gaussian beta mixing assumptions on the time series in X_t -- The temporal dependence within the data can therefore be left totally unspecified- a chi squared form of the test can be used as well.

- Fermanian(2004) has recently proposed essentially the same test but based on the original data rather than the PIT transformed data. There is evidence that the PIT transformed form of the test performs better.
- The problem with this test is that it requires the multivariate kernel density estimation which is cumbersome and leads to inaccuracy.
- Their second test is based on noting that if H_0 holds then the scalar random variable $W = \sum_{j=1}^{d} \left[\phi^{-1}(Z_j) \right]^2$ follows a χ_d^2 under H_0.
- This is just Breymann's test described above where they used the Anderson Darling Statistic but since W is not observed they based their test on the pseudo observations

from the PIT transforms and this affects the critical values of the A-D test in critical ways.

- Instead Fan et al propose a kernel based equivalent test.

$H_0^{''} : W \sim \chi_d^2$ so on $H_0^{''}, F_{\chi_d^2}(W)$ follows a uniform distribution and so they develop a test which is based on caomparing the kernel density of $g_W(.)$ with the uniform density where $g_W(.)$ is the density function of $F_{\chi_d^2}(W)$ where

$$\hat{g}_W(w) = \frac{1}{nh}\sum_{t=1}^{n} K_h(w, F_{\chi_d^2}(\hat{W}_t))$$

and

$$\hat{W}_t = \sum_{i=1}^{d}\left[\phi^{-1}(\hat{Z}_{jt})\right]^2$$

so their second test is given by

$$\hat{J}_n = \int_0^1 \left[\hat{g}_W(w) - 1\right]^2 dw$$

Again they show under beta mixing conditions that this test is asymptotically Gaussian.

- Monte Carlo with DGP

$$(U_{1t}, .., U_{dt}) \sim iid\,(1-p)C_{normal}(\rho) + pC_{t_4}(\rho)$$

Cox Tests of Non-Nested Hypotheses:
Hafez and Salmon (2004)

Comparison between Copula is a case of Non-nested testing.

Use Simulated Cox tests

$$H_f : F_\vartheta = \{ f(y \mid x, \vartheta) \}$$

$$H_g : F_\lambda = \{ g(y \mid x, \lambda) \}$$

$$T_f = l_f(\hat{\vartheta}) - l_g(\hat{\lambda}) - E_f\left[l_f(\hat{\vartheta}) - l_g(\hat{\lambda}) \right]$$

Copula Quantile Regression and Tail area Dependence

Bouyé Salmon Paper

Reading

Bouyé E. and M. Salmon; (2003) Dynamic Copula Quantile Regression and Tail Area Dependence in Foreign Exchange Markets, FERC Discussion Paper http://www2.warwick.ac.uk/go/ferc/

Bouyé E.,N. Gaussel and M. Salmon; (2002), Investigating Dynamic Dependence using Copulae, FERC Discussion Paper http://www2.warwick.ac.uk/go/ferc/

Breymann W.,A.Dias and P.Embrechts, (2003), Dependence Structures for Multivariate High Frequency Data in Finance.

Chen X., Y. Fan and A. Patton, (2004), Simple tests for Models of Dependence between Multiple Financial Time Series , with applications to US Equity returns and Exchange Rates. FERC Discussion Paper http://www2.warwick.ac.uk/go/ferc/

Chen X. and Y.Fan, (2002) Estimation of Copula Based Semiparametric Time Series Models.

Fermanian J-D. and O.Scaillet, (2004) Non-parametric Estimation of Copulas for Time Series

Fermanian J-D. and O.Scaillet, (2004), Some Statistical Pitfalls in Copula Modelling for Financial Applications .

Fermanian J-D.(2004), Goodness of Fit Tests for Copulas

Genest C. ,J-F Quessy and B. Rémillard, (2004) Goodness of fit Procedures for Copula Models Based on the Probability Integral Transform.

Genest C.,K. Ghoudi and L-P Rivest, (1995), A semiparametric estimation procedure of dependence Parameters in Multivariate Families of Distributions, Biometrika,82(3), 543-552.

Genest C. and B. Werker (2001), Conditions for the Asymptotic Semiparametric Efficiency of an omnibus estimator of dependence parameters in Copula models.

Hafez P. and Salmon (work in progress), Statistical Discrimination between Copulae using simulated Cox Tests.

Hafez P. and M. Salmon , How to choose the correct Copula.

Malevergne Y. and D. Sornette, Testing the Gaussian Copula Hypothesis for Financial Asset Dependencies

Manzotti A., F.Perez and A. Quiroz, (2002), A statistic for testing the null Hypothesis of Elliptical Symmetry, Jounral of Multivariate Analysis, 81,274-285.