Time-varying copulas on equity

Dennis van der Houwen

June 30, 2014

Abstract

I apply constant and time-varying parameter variants of the normal and Symmetrized Joe-Clayton copula to AR(p)-GARCH(1,1) models of the DAX - FTSE100 , S&P500 - FTSE100 and the S&P - S&P/TSX returns of equity price indices. Using a likelihood ratio test, I find that time-varying copulas provide a significantly better model fit than copulas with constant parameters. The North-American dependence seems to be more volatile on the short term and the European dependence seems to have a structural break around 2005. The VaR models for the 99% VaR seem to perform well, but the 95% and 90% VaR slightly overestimate the risk. Time-varying copula VaR models don't perform better than constant copula VaR models, but it does outperform the benchmark model.

Supervised by Sander Barendse
Erasmus Universiteit Rotterdam

Bachelor scriptie Econometrie & Operationele Research

Contents

1 Introduction

Recently there has been an increasing amount of evidence that stock returns are more dependent during extreme negative events than during extreme positive events (Longin and Solnik (2001) and Ang and Chen (2002))). A commonly used measure of dependence is Pearson's correlation coefficient. However, Pearson's correlation is not able to fully capture the dependence as joint asset returns are non-elliptical, Chicheportiche and Bouchaud (2012) shows this for stock returns. Pearson's correlation is only able to capture the dependency relationship when the joint distribution is elliptical. Pearson's correlation is only a linear measure of dependence and it is not able to capture non linear dependencies. Linear correlation only measures the degree of dependence and not the structure of dependence (Rachev (2009) and Necula (2010)).

Sklar (1959) introduced that every joint distribution function can be decomposed into its marginal distributions and a copula, which completely describes the dependency between the variables. Patton (2006) extended Sklar's Theorem by introducing conditional copulas and thus allowing to use copula theory in the analysis of time-varying conditional dependence.

In this paper I apply time-varying copulas on North-American and European equity returns. The Value-at-Risk (VaR) is calculated by simulation and the VaR models is evaluated with a test suggested by Christoffersen (1998) . Next to the VaR I also look at the differences between the North-American and European equity returns to see if there are any differences in dependence between the North-American equities and European equities. I'm going to test three different hypotheses:

- The time-varying VaR models calculated by Monte Carlo simulation perform better compared to constant copulas or other benchmark models.
- The dependence between North-American and European equity returns is different.
- Time-varying copulas provide a better model fit than constant copulas.

I proceed as follows. In Section 2 the data set is discussed, the methodology is discussed in Section 3. The results are discussed in Section 4 and I conclude in Section 5.

2 Data

The data used in this paper is the daily price index from two North-American and two European equity price indices. The North-American equities are the S&P 500 and S&P/TSX Composite Index. The S&P 500 is a market index which consists 500 large companies listed on the New York Stock Exchange (NYSE) or on the NASDAQ stock market. The S&P/TSX Composite Index is a market index which contains the largest companies on the Toronto Stock Exchange(TSX).

The European equities are the FTSE 100 and the Deutscher Aktien Index (DAX). The FTSE 100 is a index which contains the 100 companies listed on the London Stock Exchange with the highest market capitalization. The DAX is a index containing the 30 biggest companies trading on the Frankfurt Stock Exchange.

The daily data is obtained from Datastream and the sample is from January 1, 1990 to June 10th 2014. This sample contains 6,377 observations and all prices are converted to the US dollar. Figure 1 shows the price index of the four equities.

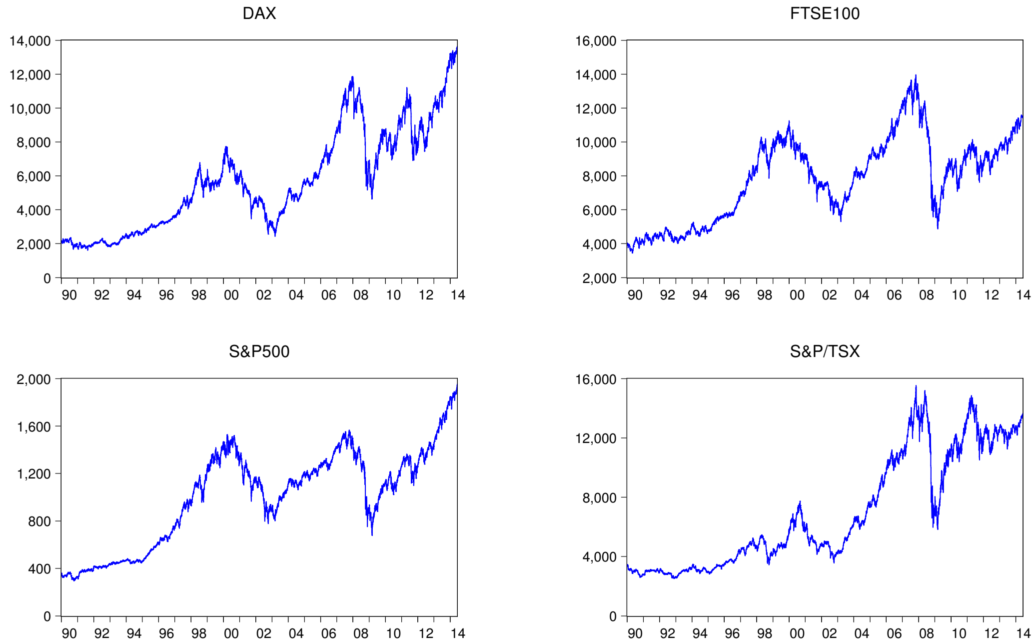

Figure 1: Plots of the complete data set of the price index for all four equities

3 Methodology

In this paper a bivariate copula is applied, the joint distribution is given by:

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)) \tag{1}$$

The copula function is given by C. $F_1(x_1)$ and $F_2(x_2)$ are the marginal distributions which are uniformly distributed.

An interesting form of dependence is the tail dependence, which is a measure of dependence in the lower or upper tail of the bivariate distribution. Basicaly lower tail dependence is defined as the limiting probably that, given an extremely small value of v, the variable u also takes an extremele small value, and vice versa for upper tail dependence. The formal definition is:

$$\tau^L = \lim_{\varepsilon \downarrow 0} \mathbb{P}[U < \varepsilon | V < \varepsilon] = \lim_{\varepsilon \downarrow 0} \frac{C(\varepsilon, \varepsilon)}{\varepsilon} \tag{2}$$

$$\tau^U = \lim_{\varepsilon \uparrow 1} \mathbb{P}[U > \varepsilon | V > \varepsilon] = \lim_{\varepsilon \uparrow 1} \frac{1 - 2\varepsilon + C(\varepsilon, \varepsilon)}{1 - \varepsilon} \tag{3}$$

Patton (2006) applied Sklar (1959) Theorem to introduce time-varying conditional copulas. This paper applies two of these time-varying conditional copulas, the Symmetrized-Joe-Clayton (SJC), which is a modification of the BB7 copula of Joe (1997). The SJC copula is given by:

$$C_{SJC}(u, v \mid \tau^U, \tau^L) = 0.5 \times (C_{JC}(u, v \mid \tau^U, \tau^L) + C_{JC}(1 - u, 1 - v \mid \tau^U, \tau^L) + u + v - 1) \tag{4}$$

Here the Joe-Clayton copula is as follow:

$$C_{JC}(u, v|\tau^U, \tau^L) = (1 - ([1 - (1 - u)^\kappa]^{-\gamma} + [1 - (1 - v)^\kappa]^{-1/\gamma} - 1)^{-1/\gamma})^{1/\kappa} \quad (5)$$

The conditioning variables τ^U and τ^L of this conditional copula are respectively the upper and lower tail dependence coefficients. Patton (2006) uses the following update equations for these tail dependence coefficients:

$$\tau_t^U = \Lambda\left(\omega_U + \beta_L \tau_{t-1}^U + \alpha_U \frac{1}{10} \sum_{j=1}^{10} | u_{t-j} - v_{t-j} |\right) \quad (6)$$

$$\tau_t^L = \Lambda\left(\omega_L + \beta_L \tau_{t-1}^L + \alpha_L \frac{1}{10} \sum_{j=1}^{10} | u_{t-j} - v_{t-j} |\right) \quad (7)$$

Where $\Lambda(x) = (1 + e^{-x})^{-1}$ is the logistic transformation to ensure that τ^U and τ^L are in $(0, 1)$ at all times. The second copula that this paper uses is the Gaussian (normal) copula. This equation is given by:

$$C(u, v|\rho) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left[\frac{-(r^2 - 2\rho rs + s^2)}{2(1 - \rho 2)}\right] dr\ ds \quad (8)$$

Here Φ^{-1} is the inverse of the standard normal $c.d.f.$ In order to make this copula time-varying, Patton (2006) uses an evoluation equation for the correlation parameter ρ:

$$\rho_t = \tilde{\Lambda}\left(\omega_\rho + \beta_\rho \cdot \rho_{t-1} + \alpha \cdot \frac{1}{10} \sum_{j=1}^{10} \Phi^{-1}(u_{t-j}) \cdot \Phi^{-1}(v_{t-j})\right) \quad (9)$$

The correlation has to stay within (-1,1) so again a logistic transformation is used $\tilde{\Lambda}(x) = (1 - e^{-x})(1 + e^{-x})^{-1}$.

For a more in depth introduction to copulas see the appendix.

As stated earlier two uniformly distributed variables are used in the copula, however since the true distribution of the marginal models is not known it's hard to find a good model which results in uniformly distributed variables. Therefore the standard residuals are first transformed into ranks. These ranks are then used for the copula models. The ranks are made as follows:

$$R^* = \frac{R_i}{n+1}, \quad S^* = \frac{S_i}{n+1} \quad (10)$$

It is computationally more demanding to use maximum likelihood to estimate the marginal model and copula parameters at the same time. Therefore pseudo maximum likelihood is used (see Genest and Favre (2007)). This means that the marginal models and copula parameters are estimated seperatly.

3.1 AR-GARCH and Data Analysis

So in this paper the daily data from four different equity price indices are used from January 1st 1990 to June 10th 2014 containing 6,377 observations. Jondeau and Rockinger (2006) suggest to eliminate the observations when a holiday occurred. When these observations are not eliminated they show the same price index as the day before and while they don't affect the dependency between stock markets during extreme events, they do affect the estimation of the return marginal distribution and subsequently the estimation of the distribution of the copula. In this paper these observations are also eliminated, this resulted in an elimination of 268 observations and now there is a total of 6109 observations. If the log differences of the returns; $r_t = 100$ x ln (P_t/P_{t-1}), where P_t is the value of the equity index at time t, is zero for atleast one of the equities, then this observation is removed for all equities.

All equities have a trend, however after taking the log differences this trend is removed. The equities are also tested on seasonality. This is done by creating four binary dummy variables and these are jointly tested to be equal. The FTSE100, S&P500 and the S&P/TSX don't show any seasonality, however the DAX does appear to have some seasonality as the wald test returns a p-value that does not reject on the 5% level. The p-value for the DAX is 0.0222, however as this is not rejected on the 1% level I won't take any measures so the models used for the equities does not change too much between each other.

This paper uses AR-GARCH models with student's t distributed error terms. First I test whether the time-series data is stationary, this is done with the Augmented Dickey Fuller test. The p-values are 0.001 , 0.000 , 0.001, 0.000 for the DAX, FTSE100, S&P500 and S&P/TSX respectively. This suggests that the null hypothesis is rejected and that the log differences are stationary. The log returns of the price indices are shown in figure 2

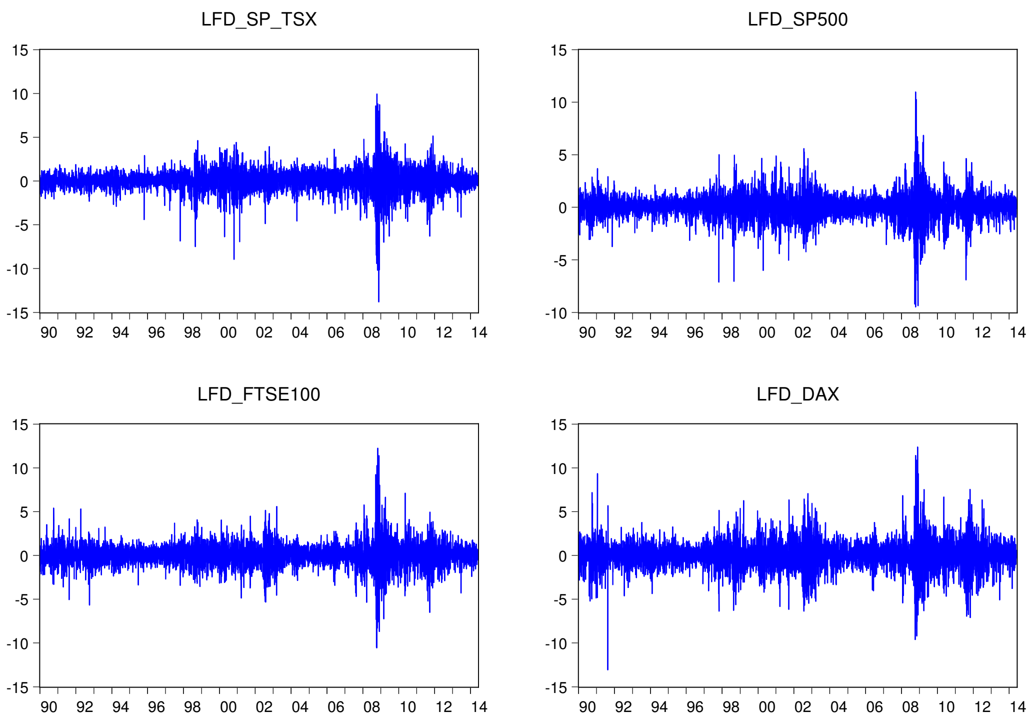

Figure 2: Plots of the complete data set for the returns of the equities

Table 1 presents some summary statistics of r_t. All equities show a negative skewness and a kurtosis that is higher than normal. The negative skewness indicates that negative returns happen more often than large positive returns. The kurtosis is higher than the kurtosis of the normal distribution, this suggests that the equities have a high peak and fat tails, so it might indeed be better to use student's t distributed error terms in the AR-GARCH models. The Q-Q-plots in figure 3 shows that the equities might not be normal. The Jarque-Bera also rejects the null-hypothesis of normality, so none of the equity price indices are normally distributed.

	DAX	FSTE100	S&P 500	S&P/TSX
Mean	0.025	0.014	0.026	0.017
Maximum	12.370	12.219	10.957	9.925
Minumum	-13.058	-10.538	-9.470	-13.789
Std. Dev.	1.560	1.272	1.152	1.227
Skewness	-0.090	-0.075	-0.240	-0.820
Kurtosis	8.240	11.617	11.658	14.893
Jarque-Bera	6995*	18903*	19134*	36685*
Number of obs.	6108			

* rejection of the null hypothesis at the 0.05 level

Table 1: Descriptive statistics for r_t of the equity price index

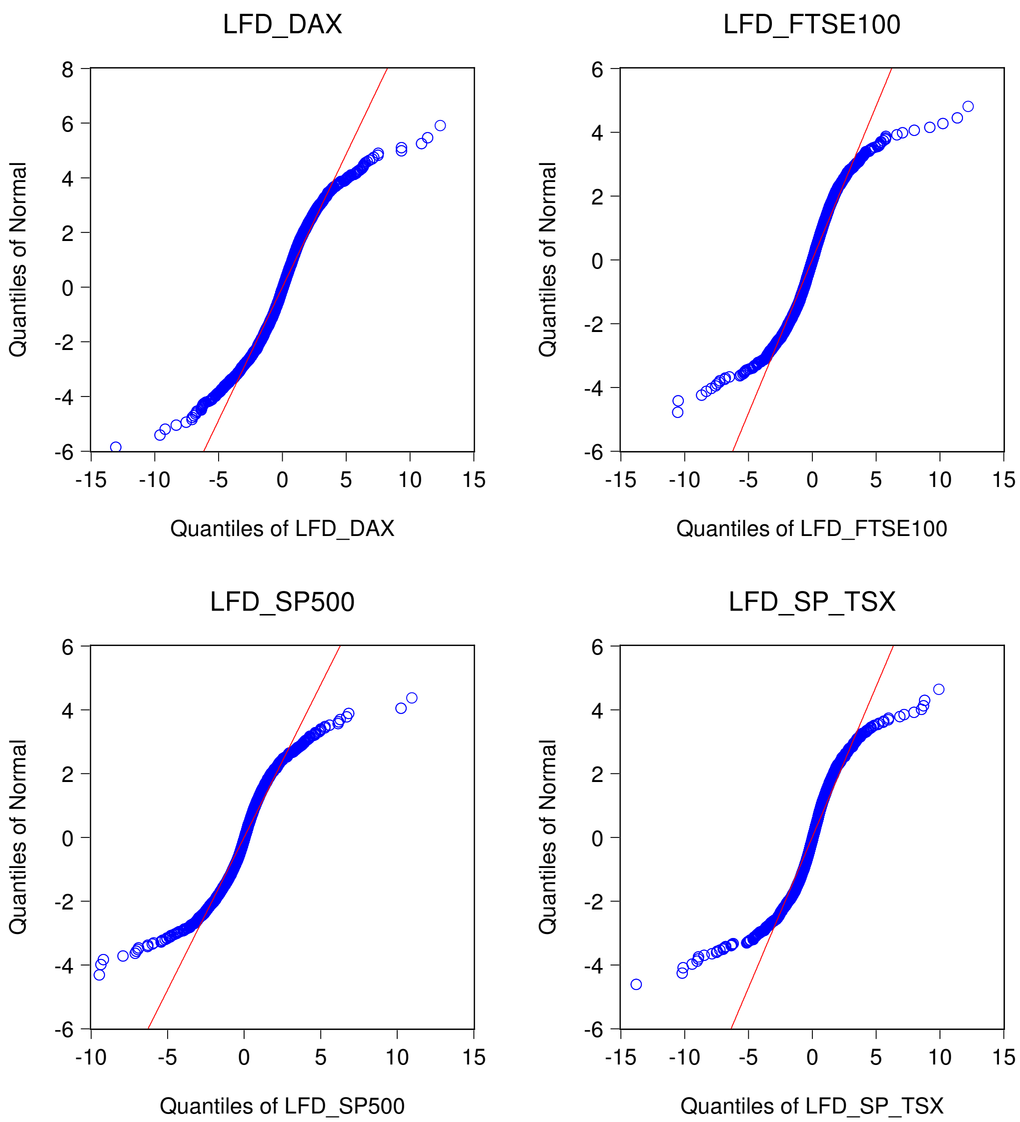

Figure 3: QQ-plots of the returns of the equity price index versus normal density

The correlation between the four equities is given in table 2. The DAX - FTSE100 have the highest correlation and after that the S&P 500 - S&P/TSX. This makes sense as these are the combinations for the European equity and North-American equity. Between the European and North-American equities the correlation is still fairly high (between 0.45 and 0.60). This suggests that there is some connection between the equities of Europe and North-America.

	DAX	FTSE100	S&P 500	S&P/TSX
DAX	1	0.7439	0.4998	0.5485
FTSE100	0.7439	1	0.4684	0.5785
S&P 500	0.4998	0.4684	1	0.6650
S&P/TSX	0.5485	0.5785	0.6650	1

Table 2: The correlation between the four returns of the equity price index.

A basic AR-GARCH(1,1) model is used for each equity (see equations 11, 12 and 13), as it is the most common model to describe financial time series (Diebold et al. (1998)). The amount of AR lag terms used is determined by the correlogram of the returns. For all returns it shows that there is no autocorrelation on a 0.05 level, so only the first AR term is added so the first expression is not just a constant (equation 11).

$$X_{i,t} = \mu_{x_i} + \sum_{j=1}^{n} \phi_{j,x_i} X_{i,t-1} + \varepsilon_{x_i,t} \tag{11}$$

$$\sigma^2_{x_i,t} = \omega_{x_i} + \beta_{x_i}\sigma^2_{x_i,t-1} + \alpha_{x_i}\varepsilon^2_{x_i,t-1} \tag{12}$$

$$\sqrt{\frac{\upsilon_{x_i}}{\sigma^2_{x_i,t}(\upsilon_{x_i}-2)}} \cdot \varepsilon_{x_i,t} \sim iid\ t_{\upsilon_{x_i}} \tag{13}$$

3.2 Value-at-Risk

A popular method used in financial markets is the so-called 'Value-at-Risk' (VaR). The VaR measures the potential loss of an asset. For example if you know the 99% VaR for 1-day-ahead for the next 100 days, then over 100 days there is on average 1 day in which you lose more than the value given by the VaR. The $VaR_t(1-q,h)$ is the q^{th} quantile of the distribution of the h-day return $r_{t+h,h}$:

$$P[r_{t+h,h} \leq VaR_t(1-q,h)] = q \tag{14}$$

So for the 1-day-ahead 95% this means that the probability that the actual return is smaller or equal to the VaR value is equal to 0.05.

The VaR is estimated by using Monte Carlo simulations. To get the 1-day-ahead VaR value for time t the data of time $t-1$ is used to estimate the GARCH model. The standard residuals from the GARCH models of the equities are used to estimate the copula model. After this two independent random vectors u_1 and v are created and put in the inverse copula function to create a new vector u_2, so $u_2 = C^{-1}_{2|1}(v|u_1)$. Now we have two vectors u_1 and u_2 that we pair up. The dependence of these two vectors are now similar to the dependence explained by the copula model, so for example if the normal copula has a correlation of 0.5 , then the vector pair should also have a correlation of 0.5. Once these vector pairs are created they are converted to the simulated standard errors by using the student's t inverse distribution with degrees of freedom equal to the ones estimated by the AR-GARCH model. The student's t distribution is used as the AR-GARCH model is estimated using the student's t error distribution. After this the simulated standard errors are put into the AR-GARCH model to get the simulated returns of the equities. For this Monte Carlo simulation the vectors contain thousand variables for each time t. A equally weighted portfolio is created from the simulated return pairs and for example the 95% VaR is then the 50^{th} smallest return. For the estimation of the AR-GARCH model a moving window is used containing the previous 500 days. 500 days are taken as this is a good datasize to estimate a GARCH model with (Hwang and Pereira (2004)).

3.3 Christoffersen test

In order to evaluate the estimated VaR values the Christoffersen test is used (see Christoffersen (1998)). The Christoffersen test contains two different likelihood ratio (LR) tests and then a third (LR) test which is basically a combination of the first two. The first test tests the hypothesis of $E[I_t] = p$ against the alternative $E[I_t] \neq p$ given independence. Here I_t is an indicator function which is equal to one if the true value is bigger than the VaR value and zero if the true value is smaller than the VaR value and p is the percent for which the VaR is evaluated. This first test is also called the unconditional coverage (uc) test. The second test tests for independence and is called the independence (ind) test. The null hypothesis is that there is no dependence against the alternative that there is dependence. A good VaR estimate should not exceed the true value for a given percent of the cases. Also there should not be any clusters of exceedings, so the VaR values should be independent of each other. The third test combines the first and the second test, so that the test looks at both of these aspects. This way it is possible to test the VaR values for both unconditional coverage and independence at the same time.

For a more in-depth explanation of the Christoffersen test, see the appendix.

3.4 Likelihood Ratio test

In order to test whether the time-varying copulas provide a better model fit than the copulas with constant parameter a likelihood ratio test is performed. The underlying marginal distributions are the same for both copulas, so it is possible to use a likelihood ratio test on the copula likelihoods. Setting parameters α and ω to zero and β to one in the equation for the time-varying copulas (equation 6, 7 and 9) is the same as the copula with constant parameters. The model under the null hypothesis with constant parameters is thus nested in the model under the alternative hypothesis with time-varying parameters.

4 Results

4.1 The Models for the Marginal Distribution

As explained in section 3.1, AR-GARCH(1,1) models are used for each equity. The results of the parameters of these marginal distributions is given in table 3. Everything except for the AR(1) values seem to be significantly different from zero. The AR(1) term for the DAX, FTSE100 and S&P500 are not significantly different from zero. However, as stated before the AR(1) terms are kept in the model so the first part is not only the constant parameter. All constant parameters are positively significant from zero, so all equity price indices increase over team. In all four cases the sum of the lagged e^2 and lagged variance is smaller than one, this suggests that the GARCH model is stationary.

	DAX	FTSE100	S&P500	S&P/TSX
Constant	0.073	0.044	0.063	0.048
	(0.014)	(0.012)	(0.010)	(0.012)
AR(1)	-0.025	-0.002	-0.022	0.123
	(0.013)	(0.013)	(0.014)	(0.013)
GARCH constant	0.018	0.014	0.007	0.006
	(0.004)	(0.003)	(0.002)	(0.002)
Lagged e^2	0.075	0.069	0.067	0.067
	(0.007)	(0.006)	(0.006)	(0.006)
Lagged variance	0.919	0.922	0.929	0.929
	(0.007)	(0.007)	(0.006)	(0.006)
Degrees of freedom	8.762	10.468	6.778	8.516
	(0.713)	(1.242)	(0.577)	(0.722)

Table 3: Results for the AR(1)-GARCH(1,1) model estimations, with the standard errors in parentheses

4.2 Results for the Copula Models

The results for the estimated normal and SJC for both the constant and time-varying copula is given by table 4. All parameters for the DAX-FTSE100 in the time-varying copulas are significant, so this might indicate that the time-varying models have a better fit than the constant models. The constant lower tail dependence is higher than the constant upper tail dependence for all three equity pairs. For the time-varying tail dependence this is not always the case as can be seen in figure 7,8 and 9. Both the S&P500-FTSE100 and S&P500-S&P/TSX have some insignificant parameters for the time-varying SJC copula. This may suggest that the tail dependence moves alot which makes it harder to find a good update equation to capture this effect. For the S&P500-S&P/TSX it is mainly the upper tail dependence that fluctuates alot as the α^U is very high. For the S&P500-FTSE100 both lower and upper tail dependence seem to fluctuate a lot as both α^U and α^L are rather high.

	DAX - FTSE100	S&P500 - FTSE100	S&P500 - S&P/TSX
		Constant normal copula	
$\bar{\rho}$	0.690	0.422	0.602
	(0.006)	(0.010)	(0.007)
Copula likelihood	1970	600	1374
		Constant SJC copula	
$\bar{\tau}^U$	0.470	0.204	0.340
	(0.012)	(0.016)	(0.015)
$\bar{\tau}^L$	0.530	0.254	0.450
	(0.010)	(0.015)	(0.011)
Copula likelihood	2024	619	1396
		Time-varying normal copula	
Constant	-0.033	-0.088	2.651
	(0.011)	(0.001)	(0,100)
α	0.375	0.016	0.714
	(0.031)	(0.000)	(0.069)
β	2.196	2.332	-2.765
	(0.045)	(0.006)	(0.081)
Copula likelihood	2119	619	1419
		Time-varying SJC copula	
Constant^U	-1.968	1.944	2.027
	(0.017)	(0.359)	(1.019)
α^U	-0.657	-9.127	-11.129
	(0.083)	(1.504)	(2.830)
β^U	4.123	-5.117	-1.452*
	(0.024)	(0.289)	(1.485)
Constant^L	-1.929	0.218*	-1.918
	(0.016)	(0.943)	(0.026)
α^L	-0.464	-6.229	-0.412
	(0.068)	(2.400)	(0.094)
β^L	4.023	0.714*	3.984
	(0.019)	(1.539)	(0.028)
Copula likelihood	2375	656	1498

*The null hypothesis is not rejected. These values are not significantly different from zero.

Table 4: Estimation output for the copula models, the standard errors are given in parenthesis

Figure 4 shows the correlation for both the constant and time-varying normal copula for the DAX - FTSE100 pair. The correlation seems to fluctuate between 0.6 and 0.8, but there are a few cases where the correlation drops down. I could not find a good explanation for this, as the price index for both the DAX and FTSE100 seem to have no special events during these time periods. After 1997 these big drops in correlation no longer happen.

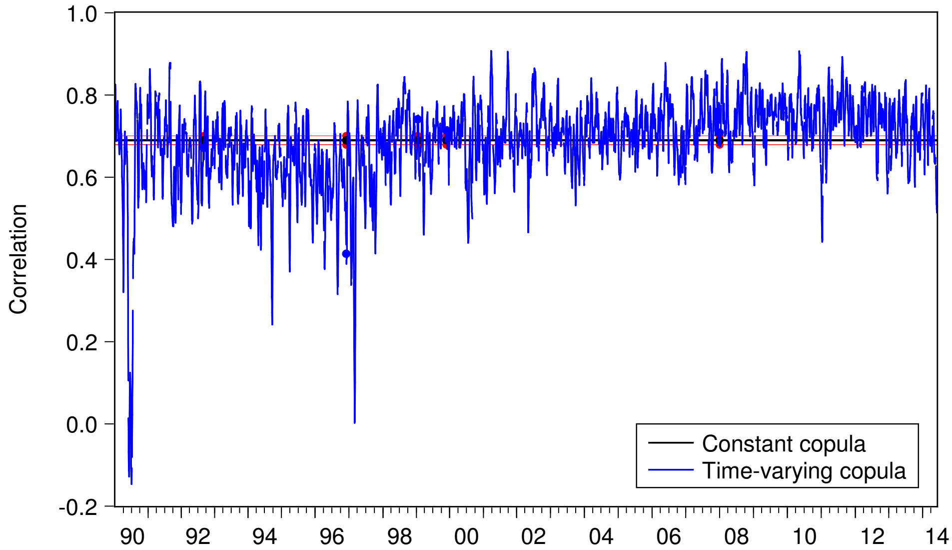

Figure 4: Correlation of the normal copula for the DAX and FTSE100

The correlation for the constant and time-varying normal copula for the S&P500 - FTSE100 pair is shown in figure 5. Here the correlation is between 0.3 and 0.5, this is lower than the DAX - FTSE100. This can be explained by the fact that the DAX and FTSE100 are European equities and have a higher correlation as shown in table 2.

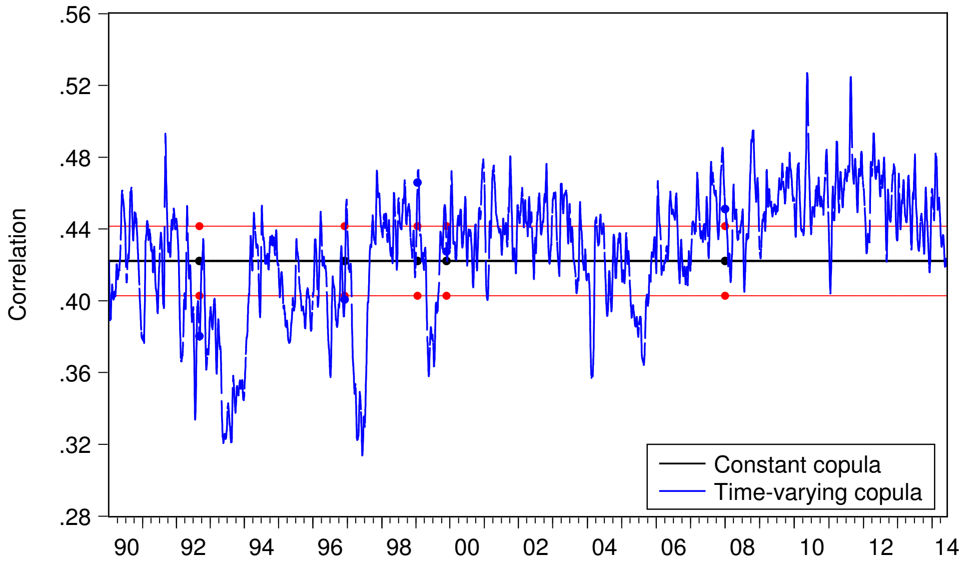

Figure 5: Correlation of the normal copula for the S&P500 and FTSE100

The correlation of the S&P500 - S&P/TSX pair is given in figure 6. The correlation of this pair is between 0.45 and 0.8 which is slightly lower than the DAX - FTSE100. This is in line with the correlation as shown in table 2. Overall this may indicate that the correlation between the European equities is slightly higher than the North-American equities, but no real conclusions can be taken from this yet.

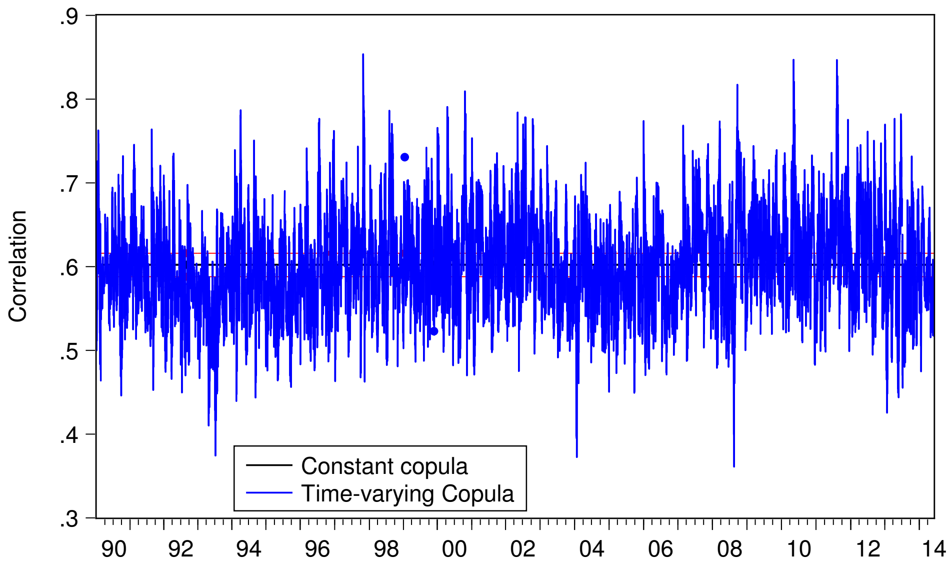

Figure 6: Correlation of the normal copula for the S&P500 and S&P/TSX

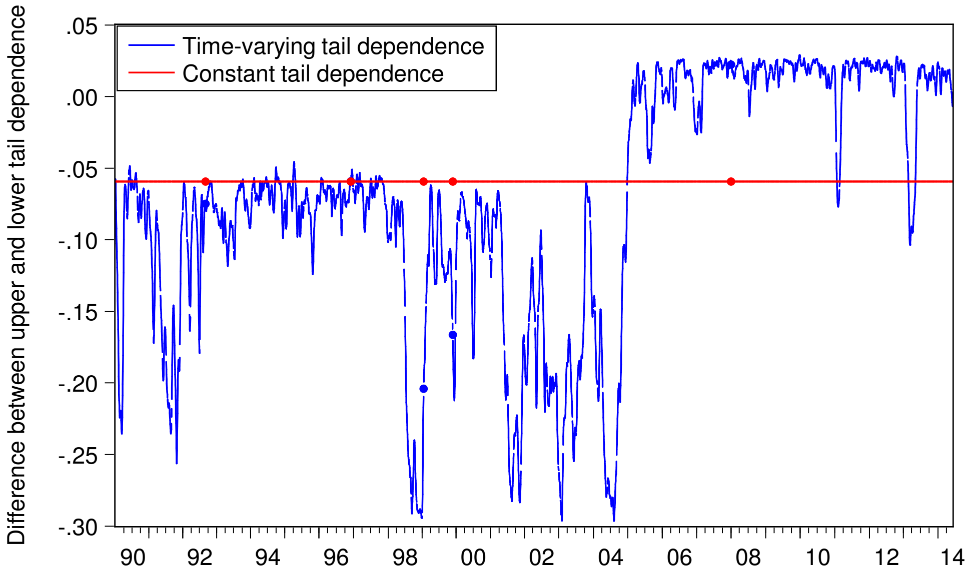

Figure 7: Difference between upper and lower tail dependence for the DAX and FTSE100

In figure 7 the difference between τ^U and τ^L is given for the constant and time-varying SJC copula of the DAX - FTSE100 pair. Up to 2005 the difference is below zero, so this suggests that joint negative events happen more often than joint positive events. After this point there seems to be some structural break, which makes the time-varying difference change from below the constant difference to above the constant difference. This means that the difference between τ^U and τ^L has become smaller.

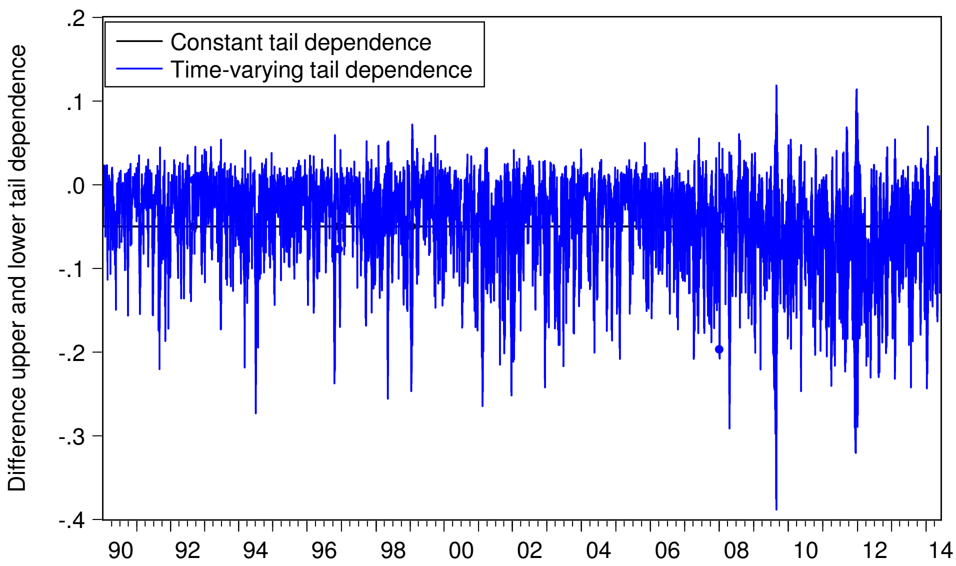

Figure 8: Difference between upper and lower tail dependence for the S&P500 and FTSE100

Figure 8 shows the difference for the S&P500 - FTSE100 pair. Just like the DAX - FTSE100 the difference is below zero, so again this suggests that there are more joint negative events than joint positive events. Around 2009 - 2012 the difference between τ^U and τ^L becomes a little more negative. This difference may be caused by the latest crisis. This suggests that during the crisis there are even more joint negative events than joint positive events.

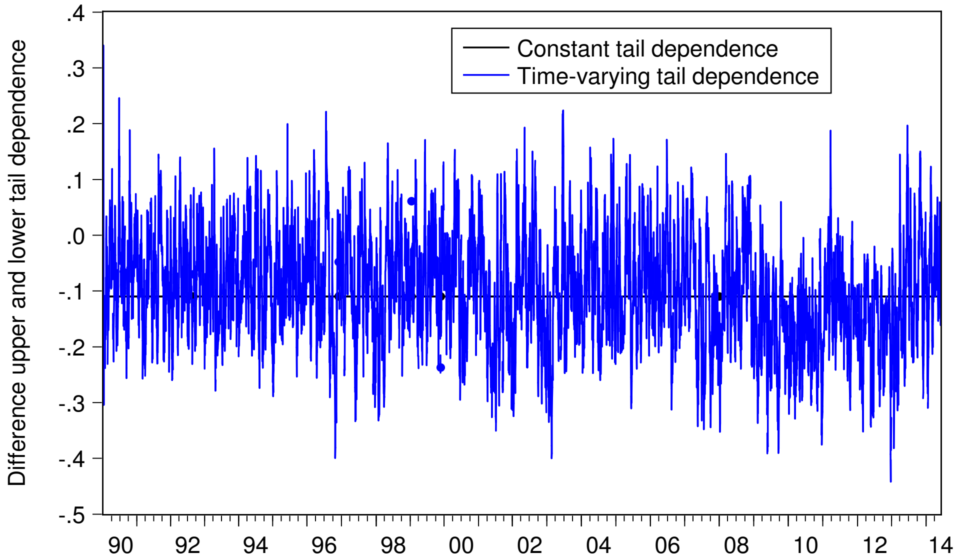

Figure 9: Difference between upper and lower tail dependence for the S&P500 and S&P/TSX

Lastly the difference between τ^U and τ^L for the constant and time-varying SJC copula of the S&P500 - S&P/TSX is given in figure 9. The difference between τ^U and τ^L seems to fluctuate more heavily than the DAX - FTSE100, so this may suggest that the upper and lower tail dependence in Europe is more steady than the North-American tail dependence on the short term. On the long term, however, the DAX - FTSE100 seems to have a structural break. This does not happen for the S&P500 - S&P/TSX. The recent crisis did seem to slightly affect the S&P500 - S&P/TSX, but there is no real change seen in the DAX - FTSE100. Perhaps this is because the crisis had the biggest impact on North-America.

4.3 Value at Risk and Christoffersen

The VaR is simulated for five different models: The SJC copula with constant parameter, SJC copula with time-varying parameter, normal copula with constant parameter, normal copula with time-varying parameter and a AR(1)-GARCH(1,1) model. The AR(1)-GARCH(1,1) model serves as a benchmark model to see how well the copula models perform. The VaR I will evaluate are the 99%, 95% and 90% 1-day-ahead VaR. Once the VaR values are simulated the Christoffersen test is applied. The results from the Christoffersen test for the pair DAX - FTSE100 are shown in table 5:

For the 99% VaR none of the combined tests is rejected, so this means the amount of hits are not significantly different from the expected hits. A total of 5.608 VaR values are estimated, so this means the expected hits for the 99%, 95%, 90% VaR are respectively 56, 280 , 561. If you compare these expected hits with the actual hits then it looks like the 99% VaR is fairly accurate, but the 95% and 90% VaR slightly overestimate the risk. The main problem seems to be the independence

		unconditional	independence	combined	#hits
99%					
	Normal constant	0.901	*	*	57
	Normal time-varying	0.901	0.612	0.873	57
	SJC constant	0.797	*	*	58
	SJC time-varying	0.901	0.612	0.873	57
	AR-GARCH	0.000	0.000	0.000	168
95%					
	Normal constant	0.058	0.000	0.000	250
	Normal time-varying	0.027	0.000	0.000	245
	SJC constant	0.077	0.000	0.000	252
	SJC time-varying	0.023	0.000	0.000	244
	AR-GARCH	0.000	0.000	0.000	492
90%					
	Normal constant	0.008	0.000	0.000	502
	Normal time-varying	0.003	0.000	0.000	496
	SJC constant	0.109	0.007	0.007	525
	SJC time-varying	0.099	0.000	0.000	524
	AR-GARCH	0.000	0.000	0.000	775

Table 5: Results for the one-day-ahead VaR for the DAX-FTSE100 Christoffersen test. The p-values of the test are given. The asterix means there is never two hits in a row, so the independence and thus the combined test can't be performed. This basically means the null-hypothesis is not rejected so there is no dependence.

test. Only the 99% VaR does not reject the independence test, the 95% and 90% do reject the independence test. This means that two hits in a row happen too often, so the hits are not independent, so it seems like there are some clusters. These clusters mainly appear when there is an increase in volatility. This makes some sense as it is harder to estimate a good VaR when the volatility increases. Compared to the benchmark AR-GARCH model the copula models seem to perform really well. In all three cases the AR-GARCH model heavily underestimates the risk which results in a large amount of hits.

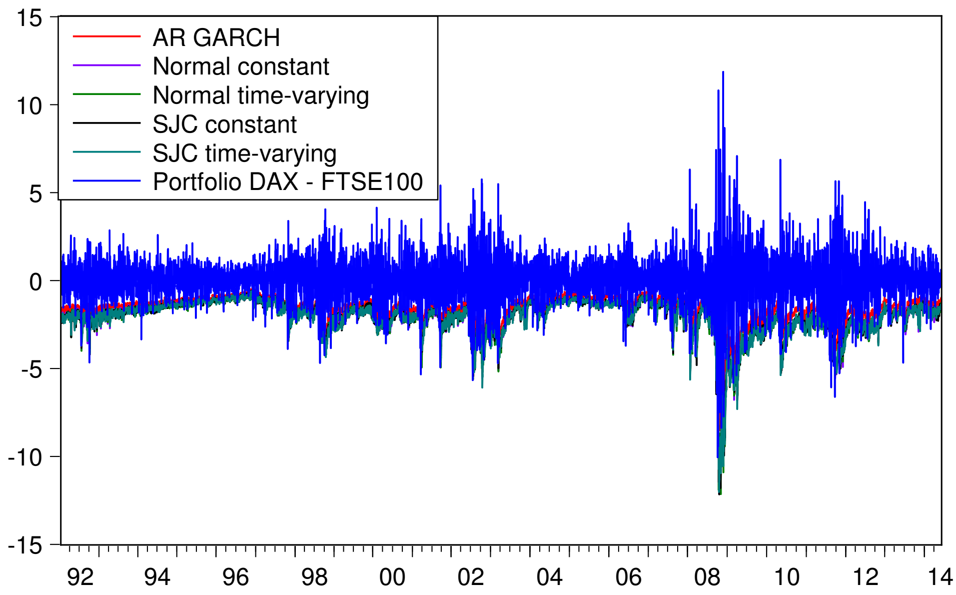

Figure 10: Plot of the 95% VaR for the DAX - FTSE100

The plot of the 95% VaR can be seen in figure 10.It is hard to see the difference between the copula models, but the AR-GARCH can be clearly seen. It looks like the AR-GARCH model does indeed underestimate the risk. In figure 11 only the last 500 observations are shown to get a better picture of how the VaR performs. Now the AR-GARCH is seen even more clearly and there are some slight differences between the copula models, but these differences are not big.

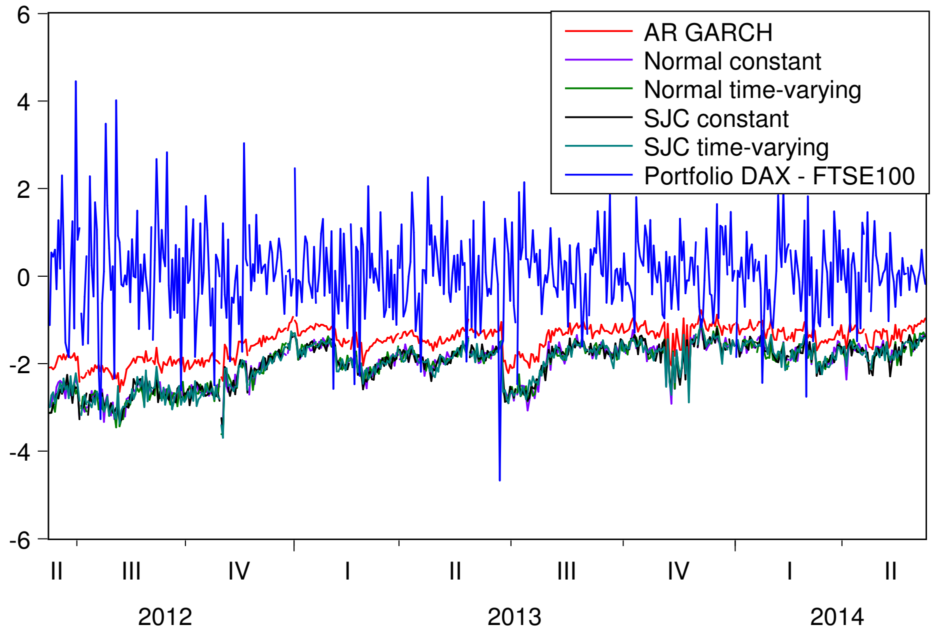

Figure 11: Plot of the 95% VaR for the DAX - FTSE100 where only the last 500 observations are shown

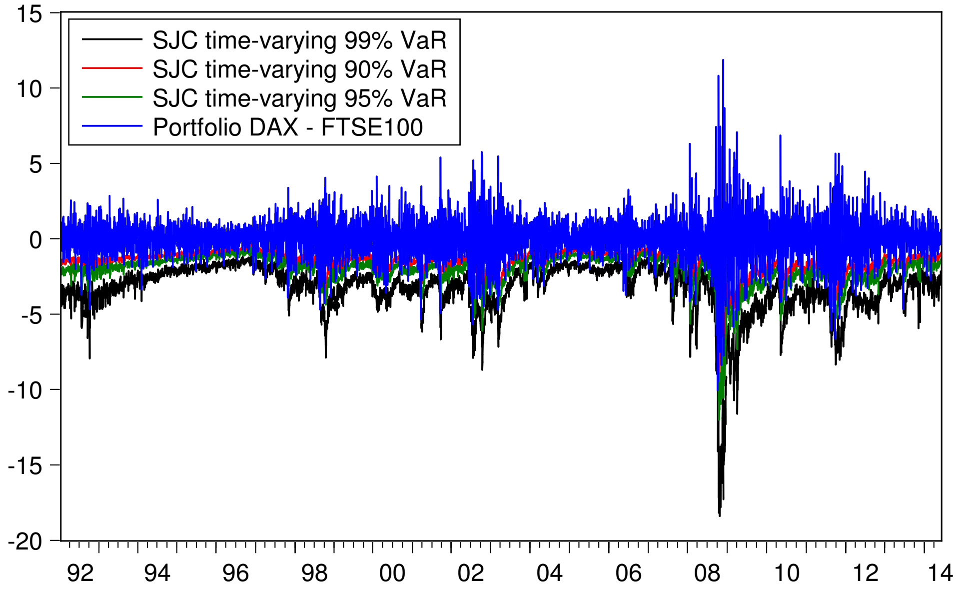

Figure 12: Plot of the DAX - FTSE100 for the 90% , 95% and 99% VaR

The results from the Christoffersen test for the S&P500 - FTSE100 are in table 6. The VaR for this portfolio seems to perform worse than the VaR for the portfolio of DAX - FTSE100. Only the normal copula with constant parameter for the 99% VaR does not reject the combined test. Even the AR-GARCH model seems to perform about equally well for the 90% VaR. Although the AR-GARCH underestimates the risk once again, for the 90% VaR this is about equal to the overestimation of risk for the copula models. The bad results of the copula models can be explained by the lower correlation between the two equities. Also this might suggest that the price indices for the North-American and European equities differ quite a bit from each other. Again these clusters happen when there is an increase in volatility.

The results from the Christoffersen test for the S&P500 - S&P/TSX are in table 7. The 99% VaR results for the copula models are not rejected, except for the SJC time-varying copula. For the 95% and 90% VaR the copula models overestimate the risk once again. The independence test is not always rejected in these cases at the 1% level, so compared to the other equity pairs this VaR model seems to be more independent. The benchmark AR-GARCH model underestimates the risk once again and performs worse than the copula models.

Overall when comparing the European and North-American pairs it looks like the European pair performs better in the unconditional test, but the North-American pair performs better in the independence test.

More plots of the VaR are located in the appendix.

4.4 Likelihood Ratio test

The likelihood ratio test described in section 3.4 is applied to the normal copula and SJC copula to see if the time-varying copula model has a better fit than the constant copula model. The p-values of the test is given in table 8. In each case the null hypothesis is heavily rejected, so the time-varying copulas provide a better fit than the constant copulas. This suggests that the correlation between the equities

		unconditional	independence	combined	#hits
99%					
	Normal constant	0.163	0.396	0.264	46
	Normal time-varying	0.034	0.309	0.062	41
	SJC constant	0.004	0.233	0.008	36
	SJC time-varying	0.006	0.024	0.002	37
	AR-GARCH	0.000	0.741	0.000	94
95%					
	Normal constant	0.003	0.005	0.000	233
	Normal time-varying	0.005	0.001	0.000	236
	SJC constant	0.001	0.070	0.001	228
	SJC time-varying	0.002	0.000	0.000	231
	AR-GARCH	0.000	0.000	0.000	370
90%					
	Normal constant	0.000	0.002	0.000	469
	Normal time-varying	0.000	0.037	0.000	476
	SJC constant	0.005	0.001	0.000	498
	SJC time-varying	0.002	0.000	0.000	493
	AR-GARCH	0.002	0.000	0.000	633

Table 6: Results for the one-day-ahead VaR for the S&P500 - FTSE100 Christoffersen test. The p-values of the test are given.

change over time. This however can't be seen in the VaR results. Maybe the moving window used for the simulation of the VaR results already compensates for the change of the correlation over time.

	p-value	
	Normal Copula	SJC Copula
DAX - FTSE100	0.000	0.000
S&P500 - FTSE100	0.000	0.000
S&P500 - S&P/TSX	0.000	0.000

Table 8: Likelihood Ratio test for constant copulas versus time-varying copulas.

		unconditional	independence	combined	#hits
99%					
	Normal constant	0.678	0.530	0.753	53
	Normal time-varying	0.885	*	*	55
	SJC constant	0.124	*	*	45
	SJC time-varying	0.006	*	*	37
	AR-GARCH	0.000	0.098	0.000	127
95%					
	Normal constant	0.037	0.071	0.022	247
	Normal time-varying	0.014	0.005	0.001	241
	SJC constant	0.023	0.183	0.031	244
	SJC time-varying	0.014	0.023	0.004	241
	AR-GARCH	0.000	0.010	0.000	428
90%					
	Normal constant	0.000	0.007	0.000	471
	Normal time-varying	0.000	0.006	0.000	464
	SJC constant	0.001	0.029	0.000	488
	SJC time-varying	0.000	0.023	0.000	480
	AR-GARCH	0.000	0.002	0.000	698

Table 7: Results for the one-day-ahead VaR for the S&P500 - S&P/TSX Christoffersen test. The p-values of the test are given. The asterix means there is never two hits in a row, so the independence and thus the combined test can't be performed. This basically means the null-hypothesis is not rejected so there is no dependence.

5 Conclusion

In this paper time-varying conditional copulas are used, suggested by Patton (2006). In total there are four different copulas used, the normal and SJC copula, both with and without time-varying parameters. These copulas are applied on equity price indices for the S&P500 , DAX , FTSE100 and the S&P/TSX.

I test three different hypotheses. The first and main hypothesis is that time-varying VaR models calculated by Monte Carlo simulation perform better compared to constant copulas or other benchmark models. The second hypothesis is that the dependence between North-American and European equity returns is different. The last hypothesis is that time-varying copulas provide a better model fit than constant copulas.

An AR(1)-GARCH(1-1) with t-distributed errors for the marginal models of the returns of the equity price indices is used.

A likelihood ratio test is used to formally test whether the copula likelihood of the time-varying model is significantly better than the copula likelihood of the copula with constant parameter. For all three pairs I find a significant rejection of the null hypothesis that the models perform equally well.

It is hard to find real differences between the dependence of the North-American and European returns. The correlation of the European returns is slightly higher than the North-American returns. The North-American dependence seems more volatile on the short term, but on the long term the European returns seem to have some structural break around 2005.

The Christoffersen test is used to test the performance of the VaR models. The 99% VaR seems fairly accurate, but the 95% and 90% VaR are often rejected by both all tests. This suggests that the VaR models are not independent. The benchmark AR(1)-GARCH(1,1) with t-distributed errors seems to perform a lot worse, so in the end the time-varying copula does not seem to perform better (or worse) than the constant copula, but it does perform better than the benchmark model.

6 Appendix

6.1 Introduction to Copulas

The use of the correlation coefficient has a few limitations. First of all empirical research in finance shows that joint asset returns are non-elliptical, see for example Chicheportiche and Bouchaud (2012) for the stock returns. Besides, it is also only a linear measure of dependence and therefore not possible to capture non linear dependencies. For example, consider a standard normal variable X. Then the correlation between X and X^2 is approximately zero even though their dependence is perfectly quadratic (Alexander (2008)). Furthermore, the correlation coefficient is not invariant under (monotonic) transformations of the variables of interest (Embrechts et al. (2002)). For example, the correlation between x and y is not the same as the correlation between $\ln(x)$ and $\ln(y)$.

Sklar (1959) shows that the dependence between random variables can be better understood through the use of a copula. We state his result for the bivariate case as we only analyze a portfolio with $n = 2$ assets in this paper.

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)), \tag{15}$$

This means that a joint distribution with two variables can be split up into two marginal distributions and a copula that governs the dependence structure. This completely describes the dependence between variables and thus is not just a linear measure such as the correlation coefficient. For example, when x and y are independent, the copula function is given by:

$$C(F_1(x_1), F_2(x_2)) = F_1(x_1)F_2(x_2) \tag{16}$$

An expression of the bivariate density function is obtained by simply differentiating with respect to each variable.

$$f(x_1, x_2) = f_1(x_1)f_2(x_2)c(F_1(x_1), F_2(x_2)), \tag{17}$$

where,

$$c(F_1(x_1), F_2(x_2)) = \frac{\delta^2 C(F_1(x_1), F_2(x_2))}{\delta F_1(x_1)\delta F_2(x_2)} \tag{18}$$

Equation (18) is also called the copula density. According to Alexander (2008), the marginal distributions $F_1(x_1)$ and $F_2(x_2)$ are uniformly distributed. When we substitute $u_1 = F_1(x_1)$ and $u_2 = F_1(x_2)$, also known as the probability integral transforms into the copula we get the copula density $c(u_1, u_2)$.

A great benefit of using copulas is that they can be applied to any marginal distribution. For example, it is possible to apply a normal copula to a marginal Student's t-distribution with v_1 degrees of freedom (df) and to another marginal Student t-distribution with v_2 df. The same copula can be applied to two completely different marginal distributions, resulting in a new joint distribution. Thus, the use of copula modeling allows us to specify a vast number of joint distributions.

An intesting form of dependence to take a closer look at is the tail dependence, which is a measure of the dependence in the lower or upper tail of a bivariate distribution. Loosely speaking, lower tail dependence is defined as the limiting probability that, given an extremely small value of v, the variable u also takes an extremely small value, and vice versa for upper tail dependence. Various copulas exhibit different forms of tail dependence. The formal defnition is given below.

$$\tau^L = \lim_{\varepsilon \downarrow 0} \mathbb{P}[U < \varepsilon | V < \varepsilon] = \lim_{\varepsilon \downarrow 0} \frac{C(\varepsilon, \varepsilon)}{\varepsilon} \tag{19}$$

$$\tau^U = \lim_{\varepsilon \uparrow 1} \mathbb{P}[U > \varepsilon | V > \varepsilon] = \lim_{\varepsilon \uparrow 1} \frac{1 - 2\varepsilon + C(\varepsilon, \varepsilon)}{1 - \varepsilon} \tag{20}$$

The standard copulas are the Gaussian copulas, which include the normal and Student's t copula. The normal copula is symmetric, meaning that $C(u_1, u_2) = C(u_2, u_1)$, so it is not able to model asymmetric dependence between two assets The Students t copula is symmetric as well however the t copula tails contain more density than the normal copula, whereas the normal copula has more mass around its center. The copula density of the normal copula is given below.

$$c(u_1, u_2; \rho) = (1 - \rho^2)^{-1/2} \exp\left(-\frac{\rho^2 \xi_1^2 - 2\rho\xi_1\xi_2 + \rho 2\xi_2^2}{2(1 - \rho^2)}\right) \tag{21}$$

Where $\xi_1 = \Phi^{-1}(u_1)$ and $\xi_2 = \Phi^{-1}(u_2)$. We show two contour plots to further illustrate that ρ is not a sufficient measure of dependence. The first figure shows the contour plot of a normal copula with a correlation coefficient of 0.5. The second plot shows the contour plot of a bivariate Clayton copula (see equation 22). The parameter of this copula is calibrated such that the corerlation coefficient is 0.5 as well, but it is obvious that the underlying dependence structure is completely different. It is better suited for modeling joint negative events.

$$c(u_1, u_2) = (\alpha + 1)(u_1^{-\alpha} + u_2^{-\alpha} - 1)^{-2-(1/\alpha)} u_1^{-\alpha-1} u_2^{-\alpha-1} \tag{22}$$

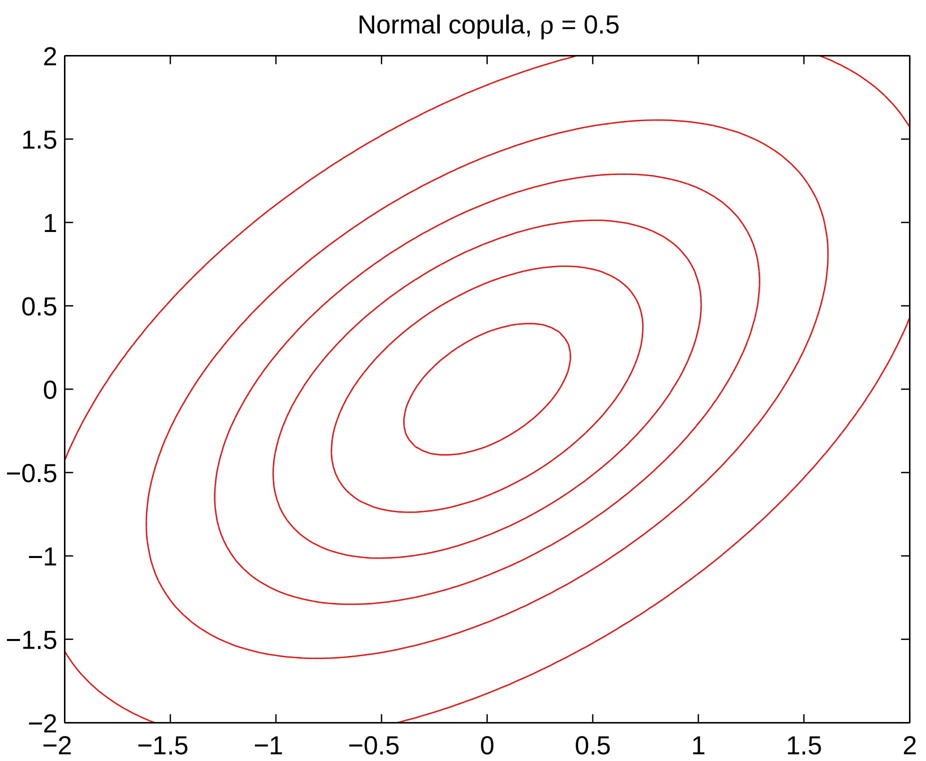

Figure 13: Bivariate normal copula contour plot

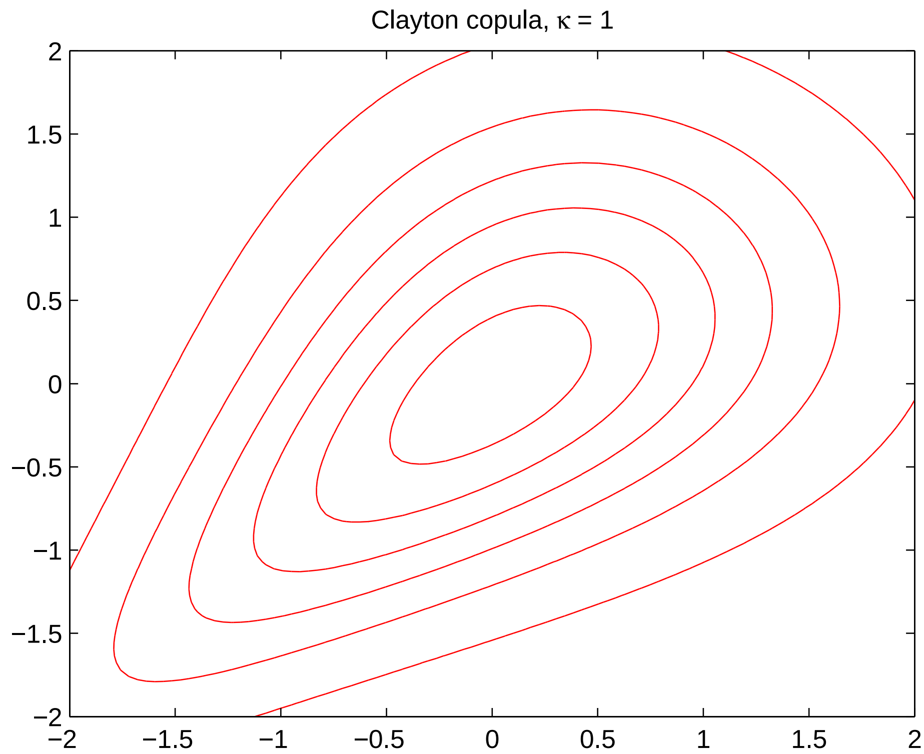

Figure 14: Bivariate Clayton contour plot

6.2 Conditional copulas

In this paper we mainly study methods and techniques introduced by Patton (2006) and apply these to the exchange rate pairs described in the data section. In his paper Patton (2006) studies asymmetric dependence between exchange rates using time-varying conditional copulas. The definition of the conditional copulas given below.

$$\begin{aligned} F_{XY|W}(x, y \mid w) =& C(F_{X|W}(x \mid w), F_{Y|W}(y \mid w) \mid w), \\ & \forall (x, y) \in \bar{\mathbb{R}} \to \bar{\mathbb{R}} \text{ and each } w \in W \end{aligned} \tag{23}$$

We follow Patton (2006) and use the symmetrized Joe Clayton copula, which is a modification of the BB7 copula from Joe (1997). The SJC copula is given below:

$$C_{SJC}(u, v \mid \tau^U, \tau^L) = 0.5 \times (C_{JC}(u, v \mid \tau^U, \tau^L) + C_{JC}(1-u, 1-v \mid \tau^U, \tau^L) + u + v - 1) \tag{24}$$

The conditioning variables τ^U and τ^L of this conditional copula respectively are the upper and lower tail dependence coefficients. We want to make τ^U and τ^L time-varying as we study the difference between copulas with time-varying parameters versus constant paramters. To do so we again follow Patton (2006), who proposes the evolution equation given in equations 25 and 26.

$$\tau_t^U = \Lambda \left(\omega_U + \beta_L \tau_{t-1}^U + \alpha_U \frac{1}{10} \sum_{j=1}^{10} \mid u_{t-j} - v_{t-j} \mid \right) \tag{25}$$

$$\tau_t^L = \Lambda\left(\omega_L + \beta_L \tau_{t-1}^L + \alpha_L \frac{1}{10}\sum_{j=1}^{10} \mid u_{t-j} - v_{t-j} \mid\right) \tag{26}$$

The evolution equations are similar to an autoregressive process and we also utilize past information of longer lags by adding the mean absolute difference of the last 10 observations. Subsequently we apply the logistic transformation $\Lambda(x) = (1+e^{-x})^{-1}$ to ensure that τ^U and τ^L stay within (0,1).

The benchmark model we use is the normal copula without time variation, see equation 27.

$$C(u, v|\rho) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left[\frac{-(r^2 - 2\rho rs + s^2)}{2(1-\rho 2)}\right] dr\ ds \tag{27}$$

Where Φ^{-1} is the inverse of the standard normal $c.d.f.$Of course, this copula has only one paramter, which is the correlation parameter ρ. We also make it time-varying by using the same evolution equation as Patton (2006).

$$\rho_t = \tilde{\Lambda}\left(\omega_\rho + \beta_\rho \cdot \rho_{t-1} + \alpha \cdot \frac{1}{10}\sum_{j=1}^{10} \Phi^{-1}(u_{t-j}) \cdot \Phi^{-1}(v_{t-j})\right) \tag{28}$$

As the correlation has to stay within (-1,1) we again apply a logistic transformation, given by $\tilde{\Lambda}(x) = (1 - e^{-x})(1 + e^{-x})^{-1}$. The update process is comparable to the SJC case, and we also look at the influence of the last 10 lags, which is done to keep the specifications comparable.

6.3 Christoffersen test

We use the VaR values to see if a copula with time-varying parameter performs better than the copula with a constant parameter. A popular test used to evaluate different VaR models is a test suggested by Christoffersen (1998). The test works as follows: Observe a sample path $\{y_t\}_{(t=1)}^T$, of the time series y_t. Define an indicator function

$$I_t = \begin{cases} 1, & \text{if } y_t \in [L_{t|t-1}(p), U_{t|t-1}(p)] \\ 0, & \text{if } y_t \notin [L_{t|t-1}(p), U_{t|t-1}(p)] \end{cases} \tag{29}$$

In the case of VaR this interval is only one-sided as we're only interested in the lower side. So for us the interval defined for the indicator function is equal to $(VaR_{(t|t-1)}(p), +\infty)$.

With the help of this indicator function Christoffersen suggests three likelihood ratio (LR) tests. First an LR test for unconditional coverage, second an LR test of independence and lastly a combination of the first two.

The first tests the hypothesis of$E[I_t] = p$ against the alternative $E[I_t] \neq p$ given independence. The likelihood under the null hypothesis is:

$$L(p; I_1, I_2, ..., I_t) = (1-p)^{n_0} p^{n_1} \tag{30}$$

And under the alternative:

$$L(\pi; I_1, I_2, ..., I_t) = (1-\pi)^{n_0} \pi^{n_1} \tag{31}$$

In the likelihood the n_0 and n_1 are the total amount of zeros and ones in the indicator, respectively.
The LR test becomes:

$$LR_{uc} = -2\log\left(\frac{L(p; I_1, I_2, ..., I_t)}{L(\hat{\pi}; I_1, I_2, ..., I_t)}\right) \sim \chi^2(s-1) = \chi^2(1), \tag{32}$$

Where $\hat{\pi} = n_1/(n_0 + n_1)$ is the maximum likelihood estimate of π, and $s = 2$ is the number of possible outcomes of the sequence (zero or one).

In this test the order of zeros and ones in the indicator vector does not matter. Only the total number of zeros and ones are important. It only tests on the coverage and doesn't hold into account any possible clusters. This is why Christoffersen uses the second LR test to test for independence.

In this second test Christoffersen tests the null hypothesis that there is independence against the alternative that there is dependence. For the alternative Christoffersen uses a binary first-order Markov chain, with transition probability matrix

$$\Pi_1 = \begin{bmatrix} 1-\pi_{01} & \pi_{01} \\ 1-\pi_{11} & \pi_{11} \end{bmatrix}, \text{where } \pi_{ij} = Pr(I_t = j \mid I_{t-1} = i). \tag{33}$$

The likelihood function for this process is:

$$L(\Pi_1; I_1, I_2, ..., I_T) = (1-\pi_{01})^{n_{00}} \pi_{01}^{n_{01}} (1-\pi_{11})^{n_{10}} \pi_{11}^{n_{11}}, \tag{34}$$

where n_{ij} is the total amount of observations with value i followed by value j.
The maximum likelihood estimation for Π_1 is:

$$\hat{\Pi}_1 = \begin{bmatrix} n_{00}/(n_{00}+n_{01}) & n_{01}/(n_{00}+n_{01}) \\ n_{10}/(n_{10}+n_{11}) & n_{11}/(n_{10}+n_{11}) \end{bmatrix} \tag{35}$$

For the null hypothesis of independence Christoffersen uses again a binary first order Markov chain,

$$\Pi_2 = \begin{bmatrix} 1-\pi_2 & \pi_2 \\ 1-\pi_2 & \pi_2 \end{bmatrix} \tag{36}$$

Here the different Markov states are independent of each other. The likelihood under the null becomes:

$$L(\Pi_2; I_1, I_2, ..., I_T) = (1-\pi_2)^{n_{00}+n_{10}} \pi_2^{n_{01}+n_{11}} \tag{37}$$

The maximum likelihood estimate is: $\hat{\pi}_2 = (n_{01}+n_{11})/(n_{00}+n_{10}+n_{01}+n_{11})$. The LR test becomes:

$$LR_{ind} = -2log\left(\frac{L(\hat{\Pi}_2; I_1, I_2, ..., I_T)}{L(\hat{\Pi}_1; I_1, I_2, ..., I_T)}\right) \sim \chi^2((s-1)^2) = \chi^2(1) \tag{38}$$

This test only tests the independence part of the time series, but it says nothing about the coverage of the series. We combine the tests for unconditional coverage and independence, so that the test looks at both of these aspects. This is the third LR test Christoffersen uses. Here Christoffersen tests the null of the unconditional coverage test against the alternative of the independence test. This gives the following LR test:

$$LR_{cc} = -2log\left[\frac{L(p; I_1, I_2, ..., I_t)}{L(\hat{\Pi}_1; I_1, I_2, ..., I_T)}\right] \sim \chi^2(s(s-1)) = \chi^2(2) \tag{39}$$

We ignore the first observation, so we simply add the first two tests together that:

$$LR_{cc} = LR_{uc} + LR_{ind} \tag{40}$$

This test allows to test time series for both unconditional coverage and for independence.

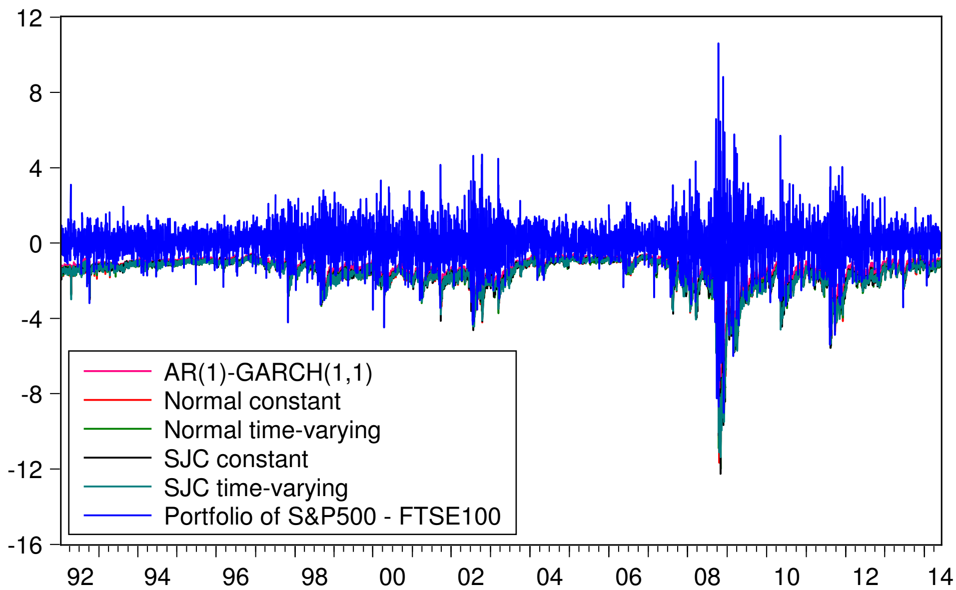

Figure 15: Plot of the 95% VaR for the S&P500 - FTSE100

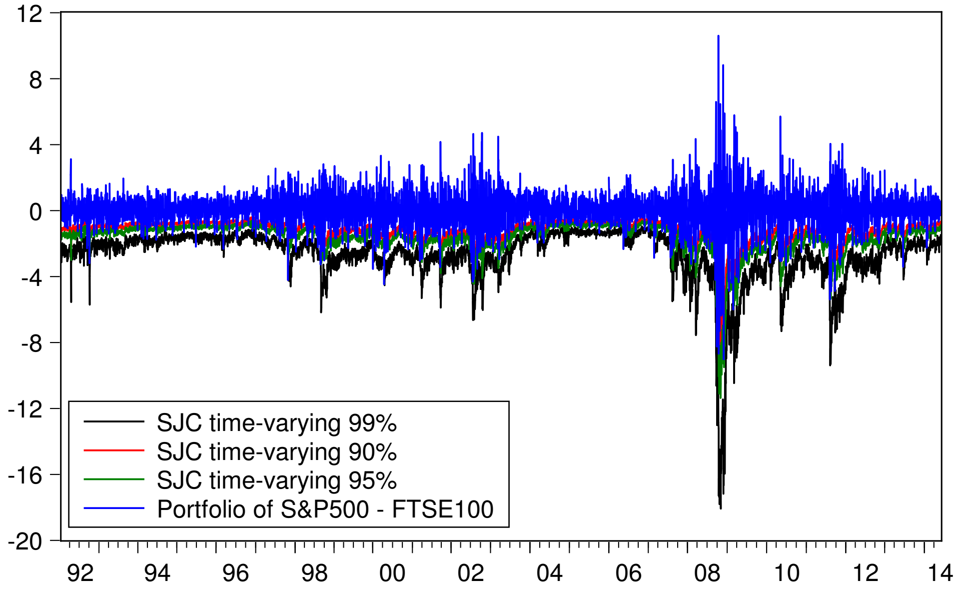

Figure 16: Plot of the SJC time-varying 99% , 95% and 90% VaR for the S&P500 - FTSE100

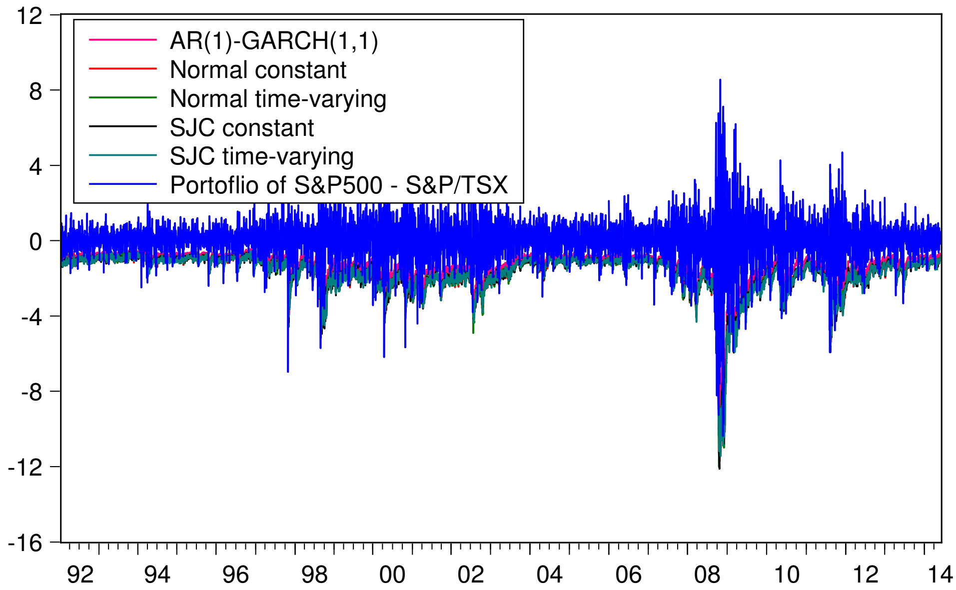

Figure 17: Plot of the 95% VaR for the S&P500 - S&P/TSX

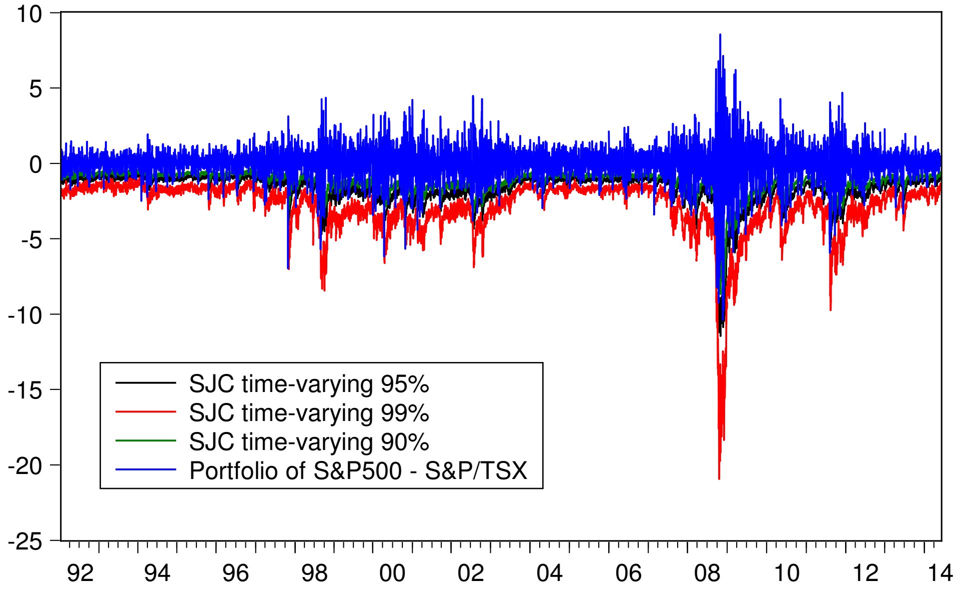

Figure 18: Plot of the SJC time-varying 99% , 95% and 90% VaR for the S&P500 - S&P/TSX

References

Carol Alexander. *Market Risk Analysis, Practical Financial Econometrics (The Wiley Finance Series) (Volume II)*. Wiley, volume ii edition, June 2008. ISBN 0470998016. URL http://www.worldcat.org/isbn/0470998016.

Andrew Ang and Joseph Chen. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443 – 494, 2002. ISSN 0304-405X. doi: http://dx.doi.org/10.1016/S0304-405X(02)00068-5. URL http://www.sciencedirect.com/science/article/pii/S0304405X02000685.

R. Chicheportiche and J. Bouchaud. The joint distribution of stock returns is not elliptical. *International Journal of Theoretical and Applied Finance*, 15(3): 1250019, 2012.

Peter F. Christoffersen. Evaluating interval forecasts. *International Economic Review*, 39(4):pp. 841–862, 1998. ISSN 00206598. URL http://www.jstor.org/stable/2527341.

Francis X. Diebold, Todd A. Gunther, and Anthony S. Tay. Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39(4):pp. 863–883, 1998. ISSN 00206598. URL http://www.jstor.org/stable/2527342.

Paul Embrechts, Alexander McNeil, and Daniel Straumann. Risk management: value at risk and beyond. 2002.

Christian Genest and Anne-Catherine Favre. Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*, pages 347–368, 2007.

Hwang and Valls Pereira. Small sample properties of garch estimates and persistence. 2004.

H. Joe. *Multivariate Models and Multivariate Dependence Concepts*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1997. ISBN 9780412073311. URL http://books.google.nl/books?id=iJbRZL2QzMAC.

Jondeau and Rockinger. The copula-garch model of conditional dependencies: An international stock market application. *Journal of International Money and Finance*, 25:827–853, 2006.

Franois Longin and Bruno Solnik. Extreme correlation of international equity markets. *The Journal of Finance*, 56(2):pp. 649–676, 2001. ISSN 00221082. URL http://www.jstor.org/stable/222577.

C. Necula. Modeling the dependency structure of stock index returns using a copula function approach. *Romanian Journal of Economic Forecasting*, 13:93–106, 2010.

Andrew J. Patton. Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2):pp. 527–556, 2006. ISSN 00206598. URL http://www.jstor.org/stable/3663514.

Michael Stein Wei Sun Rachev, S.T. Copula concept in financial markets. 2009.

A. Sklar. Fonstions de répartition á n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231, 1959.

List of Figures

List of Tables