

Increasing correlations or just fat tails?[☆]

Rachel A.J. Campbell^{a,b}, Catherine S. Forbes^c, Kees G. Koedijk^{a,b,d,*}, Paul Kofman^e

^a *Erasmus University Rotterdam, The Netherlands*

^b *Maastricht University, The Netherlands*

^c *Monash University, Australia*

^d *CEPR, United Kingdom*

^e *The University of Melbourne, Australia*

Accepted 17 January 2007

Available online 23 May 2007

Abstract

Increasing correlation during turbulent market conditions implies a reduction in portfolio diversification benefits. We investigate the robustness of recent empirical results that indicate a breakdown in the correlation structure by deriving theoretical truncated and exceedance correlations using alternative distributional assumptions. Analytical results show that the increase in conditional correlation could be a result of assuming conditional normality for the return distribution. When assuming a popular alternative distribution – the bivariate Student- t , – we find significantly less support for an increase in conditional correlation and conclude that this is due to the presence of fat tails when assuming normality in the return distribution.

© 2007 Elsevier B.V. All rights reserved.

JEL classification: G11; G14

Keywords: Exceedance correlation; Truncated correlation; Bivariate Student- t , correlation

1. Introduction

Correlation estimates are the crucial ingredient for successful portfolio management. Both market lore (e.g., Sullivan, 1995; Blyth, 1996) and academic research (e.g., Longin and Solnik, 1995, 2001; Karolyi and Stulz, 1996) suggest that there has been an increase in correlation during turbulent market conditions; resulting in a reduction in the benefits from portfolio diversification. Since it is precisely during these times that diversification is most needed, investors should be extremely concerned about a breakdown in the correlation structure.¹ Specifically investors are concerned with increasing correlation, since this is when the benefits from diversification are reduced.

[☆] The authors would like to thank participants at the Forecasting Financial Markets Conference (2000); the Q-group seminar participants in Melbourne; the IMCA seminar participants in Brisbane, Melbourne, and Sydney; the EFA 2006 meetings in Zurich; and, Stephen Figlewski, Frans de Roon, Steve Satchell, and Stefan Straetmans for valuable comments. We are especially grateful to two anonymous referees and the editor, Geert Bekaert, for many helpful comments and suggestions. All errors pertain to the authors. Kofman gratefully acknowledges support from the Australian Research Council.

* Corresponding author. RSM Erasmus University, Burg. Oudlaan 50, 3062 PA Rotterdam, The Netherlands. Tel.: +31 10 4082748; fax: +31 10 4089017.

E-mail address: ckoedijk@rsm.nl (K.G. Koedijk).

¹ Boyer et al. (1999) refer to “correlation breakdown”. Since correlation actually increases in their scenario, the breakdown in correlation leads to a reduction in the benefits from diversification.

Earlier studies found that this erosion in diversification benefits occurred more during times of greater volatility on financial markets. Longin and Solnik's (1995, 2001) findings are based on a comparison of empirical and theoretical conditional correlation measures assuming (conditional) bivariate normal distributions within an extreme value theory (EVT) framework. Ang and Bekaert (2002) found financial asset returns to be more highly correlated during bear markets but less so during bull markets. Butler and Joaquin (2002) find similar 'higher-than-normal' correlation during extreme market downturns, even when assuming fatter-tailed bivariate distributions. Unfortunately, the commonly adopted normal (or even the Student- t) distribution assumption in portfolio optimisation, implies a symmetric correlation structure incompatible with the observed empirical correlation asymmetry in bear and bull markets. Ang and Bekaert (2002), and Ang and Chen (2002) assume conditional bivariate normality, but their regime-switching (RS) model explicitly accounts for the asymmetry in conditional correlations (and volatilities).

Since increasing correlation appears to be most prevalent in (extreme) bear market conditions, a significant part of the conditional correlation literature has focussed on extreme value explanations. The increasingly popular copula class of models provides an alternative to this extreme value approach, see e.g., Embrechts, McNeil, and Straumann, (1999), Patton (2004) and Granger, Teräsvirta and Patton (2006).

However, the problem of increasing conditional correlation is not restricted to the tails of the multivariate distribution. In fact, Longin and Solnik (2001) find that normal values for empirical conditional correlations are rejected even when conditioning at very low thresholds. We therefore model the 'complete' conditional correlation structure. This approach allows us to incorporate the typical dependency structure apparent in the complete return distribution. A GARCH-type model, for example, can be used to capture time-variation in the conditional variance while simultaneously explaining a significant amount of unconditional fat-tailedness in the distribution. However, many empirical studies modelling conditional volatility of asset returns have found evidence of residual fat-tailedness in GARCH standardized returns, rejecting conditional normality.² A commonly adopted solution is then to assume a conditional Student- t distribution, see Bollerslev (1987). Longin (2005) illustrates how well this alternative model fits the empirical data. Yet, apart from Butler and Joaquin (2002), the conditional correlation literature has not yet adopted the Student- t distribution.

A number of estimators have been proposed to investigate the conditional correlation of asset returns over the complete distribution.³ We investigate the robustness of two popular conditional correlation estimators; the truncated correlation estimator, and the exceedance correlation estimator. We derive analytical expressions for these theoretical conditional correlation estimators assuming a bivariate Student- t distribution. We find that earlier results supporting increasing correlation are significantly modified when the underlying returns are jointly Student- t distributed, even after taking account of time-varying volatility in the return distribution. Our results show that the apparent breakdown in correlation structure, and the resulting reduction in diversification benefits, is largely due to the fat tailedness of financial asset return distributions.

The outline of the paper is as follows. In the following section we derive the estimation methodology for the exceedance and truncated correlation estimators when the standardized asset returns are jointly Student- t distributed. We discuss and compare the theoretical conditional correlation functions for a range of Student- t , degrees of freedom ν , including the limit case when the Student- t , converges to the normal distribution. In Section 3, we apply the truncated and exceedance correlation estimators to daily data on various international stock market index returns versus a world stock index return. We note the importance of standardizing the returns to remove the intertemporal dependency, and only then computing the empirical truncated and exceedance correlations. We compare the empirical conditional correlation estimates to their theoretical conditional correlation counterparts for the distributional assumption that best fits the standardized returns. Section 4 concludes.

2. Methodology

Conditional correlation estimates may be derived from conditioning on either the level and/or the volatility of returns of multivariate return distributions, see Barnett (1976). A multiplicity of techniques and approaches to estimate

² See for example Jansen and De Vries (1991), Huisman et al. (1998) and Jondeau and Rockinger (2003).

³ These include the multivariate GARCH models with time-varying correlations (e.g., Bollerslev et al., 1988; Karolyi, 1995 among many others).

conditional correlation highlights the absence of a unique characterization of conditional correlation. There are as many conditional correlation estimators as there are different ways of conditioning. Most of the conditioning schemes cause the conditional correlation to differ from the unconditional correlation, though not necessarily in the same direction or by the same magnitude.⁴ This makes it rather unsatisfactory for practitioners, since the conditional estimators provide alternative measures which cannot be easily compared to each other, or to their unconditional estimate. It is vital that the theoretical conditional correlation structure is used when assessing movements in the magnitude of conditional correlation estimates to their unconditional counterpart.

The two most popular conditional correlation estimators in the finance literature condition on a) a single level of return (the singly *truncated* correlation estimator) or b) on both levels of the return distribution (double truncation, better known as the *exceedance* correlation estimator). The following section briefly discusses these exceedance and truncated correlation estimators for normally distributed returns. We then develop the analogue of the exceedance and truncated correlation estimators for Student-*t* distributed returns in Section 2.2.

2.1. Truncated and exceedance correlations for the bivariate normal distribution

First, consider the case where a bear market condition applies to one ‘asset’ only – say, the S&P500 index – and we want to estimate the correlation between the US and German index returns given a large negative return on the S&P500 index. The bear market condition is characterized by a return below some threshold value, λ . The appropriate truncated correlation measure conditions on a single marginal threshold, i.e., $\text{corr}(x, y | x < \lambda)$. To assess whether an empirical estimate $\overline{\text{corr}}(x, y | x < \lambda)$ implies an increase in correlation beyond ‘normal’ correlation, we first need to compute the theoretical ‘normal’ conditional correlation for a truncated bivariate distribution.

Choose x, y to be correlated random variables driven by independent and identically distributed standard normal (SND) random variables $\varepsilon_x, \varepsilon_y$ with drift rates μ_x, μ_y , unconditional standard deviations σ_x, σ_y and unconditional correlation $\text{corr}(x, y) = \rho$, so that

$$\begin{aligned} x &= \mu_x + \sigma_x \varepsilon_x \\ y &= \mu_y + \rho \sigma_y \varepsilon_x + \sqrt{(1 - \rho^2)} \sigma_y \varepsilon_y \end{aligned} \quad (1)$$

We are interested in the correlation between (x, y) for a partitioning $Q = \{(x, y) | L < x \leq U\}$:

$$\rho_Q = \text{corr}(x, y | Q) = \frac{\text{cov}(x, y | Q)}{\sqrt{\text{var}(x | Q) \text{var}(y | Q)}}, \quad (2)$$

a ratio of truncated covariance and truncated standard deviations. Note that $\text{var}(y | Q) = \text{var}(y)$ in this singly truncated case. After some manipulation (see Johnson and Kotz, 1972 p.112) it follows that

$$\rho_Q = \text{corr}(x, y | Q) = \frac{\rho}{\sqrt{\rho^2 + (1 - \rho^2) \frac{\text{var}(x)}{\text{var}(x | Q)}}}. \quad (3)$$

We label the conditional correlation estimator in Eq. (3) as the *truncated correlation* estimator. In Eq. (3), the truncated variance of x is equivalent to the variance of a truncated normal random variable. The normal quantiles L, U define $Q = (L, U]$ so that the truncated variance of x ,

$$\text{var}(x | Q) = 1 - \left[\frac{F_N(L) - F_N(U)}{F_N(U) - F_N(L)} \right]^2 + \left[\frac{L f_N(L) - U f_N(U)}{F_N(U) - F_N(L)} \right], \quad (4)$$

⁴ A paper by Campbell et al. (2002) uses VaR-based conditional correlation estimators. Their simulation results suggest that VaR-conditional correlations do not differ from the unconditional correlation.




Fig. 1. Truncated correlation functions. a. Truncated. b. Cumulative truncated.

with $f_N(x)$ and $F_N(x)$, the probability density and distribution functions, respectively, for the standard normal distribution, further simplifies to:

$$\text{var}(x|Q) = 1 - \left[\frac{\exp(-L^2/2) - \exp(-U^2/2)}{\Pr(Q)\sqrt{2\pi}} \right]^2 + \left[\frac{L\exp(-L^2/2) - U\exp(-U^2/2)}{\Pr(Q)\sqrt{2\pi}} \right]^2, \quad (5)$$

where $\Pr(Q) = F_N(U) - F_N(L)$ is the probability of the event Q . Note that by truncating the unconditional distribution of x , the variance ratio $[\text{var}(x)/\text{var}(x|Q)]$ will exceed one, which implies that the truncated correlation will be less than the unconditional correlation.⁵

⁵ There is an exception to this rule when $Q = \{(x,y)|x \in (-\infty, U] \cup [L, \infty)\}$, a union of two partitionings as in Loretan and English's (2000b) "high-volatility" partitioning. The variance ratio will then be less than one and truncated correlation will exceed unconditional correlation. Since we want to allow for asymmetry between bear and bull market conditions, we do not consider this union of partitionings.




Fig. 2. Exceedance correlation functions.

By partitioning the bivariate distribution into intervals $(L, U]$ of equal probability $\Pr(Q)$, we obtain a U-shaped theoretical truncated correlation function with conditional correlations increasing in the tails. Fig. 1a illustrates the theoretical truncated correlation function for a bivariate normal distribution with $\rho=0.75$. For increasing truncation – decreasing $\Pr(Q)$ – truncated variance and truncated correlation will decrease monotonically (but non-linearly). Hence, the U-shaped function shifts down. Examples that use the truncated correlation estimator are Boyer et al. (1999), Loretan and English (2000a), Butler and Joaquin (2002), and Forbes and Rigobon (2002). The theoretical U-shape indicates that care needs to be taken before suggesting that tail correlations are excessive. An alternative representation is to choose a bear market interval $Q^- = (-\infty, U]$ and compute the truncated correlations for increasing quantile U . We call this the cumulative truncated correlation function. For the bivariate normal distribution, the cumulative truncated correlation function has an inverted U, or ‘tent,’ shape.⁶ Fig. 1b illustrates this theoretical cumulative truncated correlation function for a bivariate normal distribution with $\rho=0.75$.

Next, consider the correlation between asset returns x and y during bear market conditions that apply to both assets. We are now interested in the correlation between (x, y) for a partitioning $Z = \{(x, y) | L < x \leq U, L < y \leq U\}$:

$$\rho_Z = \text{corr}(x, y | Z) = \frac{\text{cov}(x, y | Z)}{\sqrt{\text{var}(x | Z) \text{var}(y | Z)}}, \quad (6)$$

a ratio of a doubly truncated covariance and two truncated standard deviations. Ang and Chen (2002, p.487–488) derive the analytical expressions for a bivariate normal distribution, which are considerably more complicated than Eq. (3). We label the conditional correlation estimator in Eq. (6) as the *exceedance correlation* estimator. Longin and Solnik (1995, 2001) are the key proponents of this estimator.

Fig. 2 illustrates the theoretical exceedance correlation function for a bivariate normal distribution with $\rho=0.75$. Similar to the cumulative truncated correlation function, the exceedance correlation function is ‘tent’ shaped. A ‘tent’ shape suggests that larger-than-normal empirical tail exceedance correlations are even more ‘excessive’ than first thought.

2.2. Truncated and exceedance correlations for the bivariate Student-t distribution

The truncated and exceedance correlation functions in Eqs. (3) and (6) allow us to identify theoretical conditional correlations for a given truncation of the bivariate distribution. We can then compare empirical truncated/exceedance

⁶ Note that the function consists of two ‘parts.’ The partitioning is accumulating from quantile 0 to quantile 0.5, i.e. $Q^- = (-\infty, U]$ with increasing U ; while the partitioning is decumulating from quantile 0.5 to quantile 1, i.e. $Q^+ = [L, \infty)$ with increasing L .

correlations to the relevant theoretical equivalents and test whether correlation is indeed excessive in the tails. Of course, to make this a valid exercise, the empirical returns will have to satisfy the assumptions underlying the truncated (co-)variances in Eqs. (3) and (6). A problem arises if the data are not bivariate normally distributed as assumed, but instead have fatter tails than bivariate normality implies. This would naturally lead to a further dispersion in the tail observations, and to an even steeper increase in truncated tail correlations. Hence, it could give the mistaken impression of excessive conditional correlation.

According to Boyer et al. (1999), as long as the bivariate density is elliptic, truncated correlation is defined as in Eq. (3). This will hold for a bivariate Student- t_r density. Of course, the truncated variance expression in Eq. (5) will change for the Student- t_r distribution. If we assume that the underlying density is jointly Student- t_r distributed, then the truncated variance becomes:

$$\text{var}(x|Q, r) = \frac{L f_{r-2}(L) - U f_{r-2}(U) + F_{r-2}(U) - F_{r-2}(L)}{\left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L\right) \right]} - \frac{[f_{r-2}(L) - f_{r-2}(U)]^2}{\left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L\right) \right]^2} \quad (7)$$

with $f_r(x)$ and $F_r(x)$, the probability density and distribution functions, respectively, for the Student- t_r distribution with r degrees of freedom. The derivation of Eq. (7), along with expressions for f_r and F_r are given in Appendix A. The expression in Eq. (7) is still computationally straightforward. In graphical terms it also implies a U-shaped truncated correlation function. Fig. 1a illustrates the theoretical truncated correlation function for a bivariate Student- t_r distribution with $\rho=0.75$, and $r=3,5,7,9,11,15$. The increased tail dispersion of the Student- t_r (in comparison with the normal) distribution generates a steeper U-shape (than for the normal). Note that for the central intervals, the Student- t_r truncated correlations are smaller than their normal equivalents. Unfortunately, significant theoretical differences between the normal and Student- t_r truncated correlations only appear in the outer tails of the distribution. These are also the regions where conditional correlations are less precisely estimated and the standard errors are largest. A more discriminating approach is to consider the cumulative truncated correlation functions. Whereas the bivariate normal assumption leads to a ‘tent’ shaped cumulative truncated correlation function, a bivariate Student- t_r assumption leads to a U-shaped cumulative truncated correlation function for sufficiently ‘small’ degrees of freedom. Fig. 1b illustrates the theoretical cumulative truncated correlation function for a bivariate Student- t_r distribution with $\rho=0.75$, and $r=3,5,7,9,11,15$. When r increases, the U-shape becomes progressively flatter and ultimately inverts when the bivariate Student- t_r converges to a bivariate normal distribution.

For exceedance correlations, we obtain a similar result. Assuming a bivariate normal density, Ang and Chen (2002) derive a ‘tent’ shaped exceedance correlation function. Assuming a bivariate Student- t_r density, on the other hand, leads to a U-shaped exceedance correlation function for a sufficiently ‘small’ r . Fig. 2 illustrates the theoretical exceedance correlation function for a bivariate Student- t_r distribution with $\rho=0.75$, and $r=3,5,7,9,11,15$. When r increases and the Student- t_r distribution converges to normality, the U-shape ultimately inverts to the normal tent shape. The derivation of the theoretical exceedance correlations assuming a bivariate Student- t_r distribution is given in Appendix B.

2.3. Testing for excess correlation

To test whether empirical conditional correlations are significantly different from the theoretical conditional correlations, we need to compute appropriate standard errors. Although calculation of the analytical solutions for these standard errors could follow in a manner similar to the expectations computed in the Appendices A and B, this type of solution would not account for the estimation error that occurs in the standardization of the raw returns to ensure that these are independently and identically distributed. To incorporate this standardization error we need to resort to simulated standard errors. These standard errors allow us to compute multiple point-wise confidence intervals around the theoretical conditional correlation functions. To avoid a possible size distortion with these point-wise tests, we also compute a joint significance test. Ang and Chen (2002) derived their H -test for this purpose, where the H statistic

$$H(F) = \sqrt{\sum_{i=1}^N w(\lambda_i) \cdot (\rho_Q(\lambda_i, F) - \bar{\rho}_Q(\lambda_i))^2}. \quad (8)$$

is a weighted average of the squared deviation between the distribution implied conditional correlations $\rho_Q(\lambda_i, F)$, and the empirically observed conditional correlations $\bar{\rho}_Q(\lambda_i)$. We choose the weights $w(\lambda_i)$ for each truncation/exceedance level λ_i proportional to the sample size used to compute the empirical truncated/exceedance correlations:

$$w(\lambda_i) = \frac{T_i}{\sum_{j=1}^N T_j} \text{ where } T_i \text{ is the sample size of each of } N \text{ partitionings } \lambda_i.$$

To distinguish bull from bear market fit, we compute sub-tests $H^+(F)$ and $H^-(F)$ from upside, respectively downside truncation/exceedance levels. And finally, we compute

$$AH(F) = \sum_{i=1}^N w(\lambda_i) (\rho_Q(\lambda_i, F) - \bar{\rho}_Q(\lambda_i)) \quad (9)$$

which may be negative if the estimated truncated/exceedance correlations are less than the distribution-implied truncated/exceedance correlations. When F is assumed to be normal, standard errors of the H statistics can be estimated by GMM using the [Newey and West procedure \(1987\)](#).⁷ When we assume F to be a Student- t_r distribution, we also have closed-form solutions for the distribution implied conditional correlations $\rho_Q(\lambda_i, t_r)$, and can use the same procedure. However, since we first estimate standardized returns to compute the conditional correlations, we need to resort to simulated standard errors for the H statistics as well. The details of our simulation procedure are given in Appendix C.

3. Empirical results

The previous section highlights the risk of too easily concluding that correlation strengthens – and diversification evaporates – during extreme market conditions. By conditioning a correlation estimate on a partition of the bivariate distribution, one should expect a different (conditional) correlation estimate than the unconditional correlation. The difference depends on the type of partitioning (exceedance, truncated or cumulative truncated), but also on the underlying bivariate distribution. When the underlying distribution is fatter tailed than normal, it seems that the difference between conditional and unconditional correlation is more pronounced. Correct inference on conditional correlation requires identification of the underlying distribution.

3.1. Data, standardized returns and preliminary statistics

We apply the exceedance and (cumulative) truncated correlation estimators to a data set that consists of daily stock market index data collected from Datastream for the USA, UK, French, and German stock indices, as well as for the MSCI All Country World Index (a free float-adjusted market capitalization index that is designed to measure equity market performance in the global developed and emerging markets). All series represent unhedged total return indices. The sample period extends from 2 January 1990 until 3 March 2005, i.e., 3958 daily observations. We note that this sample period covers bull market, bear market, and ‘normal’ market episodes. Our sample length exceeds [Longin and Solnik’s \(2001\)](#) sample since we want to address possible market cycle dependency of correlation estimates. We also use a higher sampling frequency⁸ to improve the efficiency of the estimation of the tail correlations.

We use continuously compounded daily returns on the MSCI World, S&P500, FTSE100, CAC40, and the DAX30 indices. Descriptive statistics are presented in [Table 1](#). The daily returns indicate an average annualized return on equity markets between 6 and 10% over the entire sample period. The annualized standard deviation varies between 12% and 22%. The return data certainly contains extremes with a maximum *daily* return in excess of 7% and a minimum *daily* return of less than 8%. The French (CAC40) and German (DAX30) stock markets are more volatile and generate more

⁷ Details are given in [Ang and Chen \(2002, Appendix C\)](#).

⁸ [Longin and Solnik \(2001\)](#) use monthly data.

Table 1

Summary statistics for index returns

	World	S&P500	FTSE100	CAC40	DAX30
Annualized mean return	7.10%	10.00%	8.31%	7.60%	6.12%
Annualized standard deviation	12.29%	16.00%	16.26%	21.09%	22.04%
Maximum daily return	5.16%	5.58%	5.90%	7.00%	7.27%
Minimum daily return	−4.42%	−7.11%	−5.89%	−7.68%	−8.24%
Skewness	−0.12 ^a	−0.11 ^a	−0.09 ^a	−0.06 ^a	−0.17 ^a
	(−0.17 ^a)	(−0.33 ^a)	(−0.09 ^a)	(−0.12 ^a)	(−0.16 ^a)
Kurtosis	6.07 ^a	6.92 ^a	6.15 ^a	5.69 ^a	6.27 ^a
	(4.45 ^a)	(4.91 ^a)	(3.99 ^a)	(3.73 ^a)	(4.33 ^a)
Normality test statistic ^b	780.6 ^a	1130.0 ^a	818.9 ^a	648.3 ^a	848.8 ^a
	(231.8 ^a)	(328.5 ^a)	(125.3 ^a)	(74.7 ^a)	(201.0 ^a)

The table gives the summary statistics for daily total returns for the following indices: MSCI World Index, S&P500 Composite Index, FTSE100 All Share Index, CAC40 Index, and the DAX30 Performance Index over the period 2 January 1990 – 3 March 2005 ($N=3958$ observations). Summary statistics for the standardized daily total returns are given in parentheses (.).

^a Indicates significantly different from normal distribution values at 95% confidence level.

^b The normality test is given in Doornik and Hansen (1994), which is based on the Jarque and Bera normality test and corrects for the fact that sample kurtosis approaches normality very slowly. The test is chi-squared distributed with 2 degrees of freedom.

Table 2

Filtering with asymmetric GJR-GARCH(1,1)- t_r

	MSCI World	S&P500	FTSE100	CAC40	DAX30
<i>Maximized log-likelihoods</i>					
GARCH(1,1)	14,056.7	13,073.2	13,021.0	11,890.4	11,927.0
GJR-GARCH(1,1)	14,088.2	13,110.6	13,050.2	11,923.5	11,946.6
GARCH(1,1)- t	14,116.8	13,163.6	13,047.8	11,917.5	11,975.6
GJR-GARCH(1,1)- t	14,144.6	13,186.8	13,076.5	11,947.5	11,990.3
<i>Parameter estimates for GJR-GARCH(1,1)-t_r</i>					
Asymmetry	0.0033	0.0042	0.0049	0.0068	0.0042
	(0.0005)	(0.0005)	(0.0005)	(0.0007)	(0.0008)
t -df (r)	8.17	7.08	13.18	12.09	9.77
	(0.99)	(0.86)	(2.71)	(2.13)	(1.43)

The table gives the maximized log-likelihoods and asymmetry and Student- t_r degrees of freedom parameter estimates for the full sample period ($N=3958$ observations) for the following daily index return series: MSCI World Index, S&P500 Composite Index, FTSE100 All Share Index, CAC40 Index, and the DAX30 Performance Index. Robust standard errors are given in parentheses. The results are given for the GARCH(1,1) model, the asymmetric GJR-GARCH(1,1) model (see Glosten et al., 1993), and both models with Student- t_r distributed errors.

extreme returns than the US and UK stock markets. All return series exhibit significant skewness as well as significant excess kurtosis, relative to normally distributed values for skewness and kurtosis. The normality test is therefore strongly rejected for every series.

The exceedance and (cumulative) truncated correlation estimators in Section 2 require that the univariate return observations are independently and identically distributed.⁹ Given typical empirical results, we expect that the variance of the returns is in fact time-dependent. We need to remove this univariate time-dependency to isolate the impact on conditional correlations of mistakenly assuming a normal distribution, when the true underlying distribution is in fact Student- t_r . We therefore first filter the univariate series by estimating a sequence of univariate GARCH(1,1) models and then compute the standardized residuals. Of course, this implies that we will need to account for filtering error in the computation of appropriate standard errors for the conditional correlations and H statistics.

⁹ It is possible to evaluate conditional correlations for non-independent, non-identically distributed returns, see e.g., Ang and Bekaert (2002) and Ang and Chen (2002). As there is no closed form solution for these cases, the implied conditional correlations need to be obtained by simulation. Note that we do not exclude the possibility of time-varying conditional correlations, but rather investigate the impact of such time-variation explicitly. We do, however, exclude the possibility that univariate dependency structures interact with the bivariate dependency. This is similar to e.g. Loretan and English (2000a,b) and Longin and Solnik (2001).

Table 3
Index return correlations

	MSCI World	S&P500	FTSE100	CAC40
<i>Panel A — Raw returns</i>				
S&P500	0.76 [0.86]			
FTSE100	0.59 [0.67]	0.40 [0.43]		
CAC40	0.59 [0.69]	0.41 [0.47]	0.75 [0.81]	
DAX30	0.61 [0.69]	0.42 [0.49]	0.65 [0.73]	0.76 [0.83]
<i>Panel B — Standardized returns</i>				
S&P500	0.71 [0.85]			
FTSE100	0.53 [0.64]	0.38 [0.44]		
CAC40	0.52 [0.66]	0.37 [0.45]	0.70 [0.76]	
DAX30	0.52 [0.66]	0.34 [0.47]	0.57 [0.68]	0.68 [0.79]

The table gives the unconditional correlation matrix for the full sample period 2 January 1990–3 March 2005 ($N=3958$ observations) for the following daily index return series: MSCI World Index, S&P500 Composite Index, FTSE100 All Share Index, CAC40 Index, and the DAX30 Performance Index. Numbers in square brackets are estimates for the sub-period 21 May 1997–3 March 2005 ($N=2032$ observations).

The GARCH(1,1) parameters are invariably highly significant. Table 2 provides the maximized log-likelihoods for the GARCH(1,1), the asymmetric GJR-GARCH(1,1), the GARCH(1,1) with Student- t_r disturbances, and the asymmetric GJR-GARCH(1,1) with Student- t_r disturbances. The asymmetric GJR-GARCH(1,1) uses the [Glosten et al. \(1993\)](#) specification. Successive likelihood ratio tests show that the GJR-GARCH(1,1)- t_r specification dominates the alternatives for all five series.

Table 2 also gives the estimates for the asymmetry parameter and the t -degrees of freedom parameter from the GJR-GARCH(1,1)- t_r specification. The asymmetry parameters are invariably significant indicating a larger volatility impact for negative returns. The t -degrees of freedom parameters vary from 7 to 13. The summary statistics for the standardized residuals from these GJR-GARCH(1,1)- t_r specifications (Table 1, in parentheses) indicate that accounting for asymmetry and Student- t_r disturbances absorbs much of the fat-tailedness of the raw returns with a substantial reduction in kurtosis. Note that the residual kurtosis is inversely related to the t -degrees of freedom parameter estimate. Skewness, on the other hand, does not seem to be reduced and the normality test is still rejected for each of the 5 series.

Table 4
Bivariate Student- t_r degrees of freedom and coskewness estimates

	Coskewness with MSCI world	t -degrees of freedom (r)			
		MSCI world	S&P500	FTSE100	CAC40
S&P500	−0.066* [0.079*]	7.09 [5.52]			
FTSE100	−0.110* [−0.035]	9.75 [4.87]	10.18 [4.15]		
CAC40	−0.106* [−0.003]	9.00 [4.98]	9.07 [4.35]	9.18 [4.34]	
DAX30	−0.125* [−0.075*]	8.34 [5.32]	8.58 [4.68]	9.44 [4.68]	7.58 [4.14]

The table gives estimates for the coskewness between country index returns and MSCI World index returns, and bivariate Student- t_r degrees of freedom parameter estimates for the standardized index return series, estimated by maximum likelihood. Coskewness is defined as in Harvey and Siddique (2000, p.1276). An asterisk indicates significant coskewness at the 5% level. Numbers in square brackets are estimates for the second sub-period 21 May 1997–3 March 2005.




Fig. 3. Truncated correlations for the CAC40 versus DAX30 returns.

Table 3 gives the unconditional correlation estimates, for each combination of raw and standardized index returns. The unconditional correlation estimates vary from a high of 0.71 (between standardized returns for MSCI World and S&P500) to a low of 0.34 (between the standardized S&P500 and DAX30). The results underline greater co-movement between European stock market returns and between MSCI World and US returns (partially explained by the substantial US weight in the MSCI World index).

The results that follow only use the standardized index return series. Table 4 gives the bivariate Student- t_r degrees of freedom and coskewness estimates.¹⁰ The bivariate Student- t_r degrees of freedom vary from a low of 7 (MSCI World versus S&P500) to a high of 10 (S&P500 versus FTSE). The coskewness estimates suggest the possibility of asymmetry in our conditional correlation functions with significant negative coskewness in all country index returns with the MSCI World index returns.

3.2. Truncated and exceedance correlation results

The (cumulative) truncated correlation functions and empirical estimates of the (cumulative) truncated correlations are given in Fig. 3, respectively Fig. 4 for the CAC40 versus DAX30 standardized returns. Truncated correlation estimates are conditioned on a threshold return for the DAX30. In Fig. 3, the truncated correlations are estimated for non-overlapping quantiles of 5% each. In Fig. 4, the cumulative truncated correlations are estimated for accumulating quantiles from 5% to 50% downside coverage, and for 'de'cumulating quantiles from 50 to 5% upside coverage.

Fig. 3 shows that it is nearly impossible to discriminate between the normal and Student- t_r truncated correlation models. Proper distinction only occurs in the outer tails when we choose 'narrow' intervals for the truncation¹¹ $[L, U]$. Of course, at the same time the lack of a sufficient sample size frustrates robust inference. The empirical estimates therefore fit within the confidence bands for both distributional assumptions. Fig. 4, on the other hand, shows that the competing distributional assumptions are clearly distinguished for the cumulative truncated correlation functions. The Student- t_r assumption provides a marginally better fit for the downside conditional correlations, while the normal assumption provides a better fit for the upside conditional correlations.

¹⁰ Coskewness is defined as in Harvey and Siddique's (2000, p.1276) direct measure β_{SKD} , i.e., the contribution of a country's index returns to the skewness of a broader portfolio (the MSCI World index) returns. A negative measure implies that a country index adds negative skewness to the MSCI World index returns.

¹¹ These intervals contain less than 5% of the observations.




Fig. 4. Cumulative truncated correlations for the CAC40 versus DAX30 returns.

Fig. 5 gives the theoretical exceedance correlation functions and the empirical exceedance correlation estimates for the CAC40 versus DAX30 returns. We observe a good fit for both upside and downside conditional correlations with the Student- t_r model. If we had assumed a normal distribution, we would have concluded that there is evidence of excess conditional correlation. For the more appropriate Student- t_r assumption, this is no longer the case.

3.3. Stability of the truncated and exceedance correlation results

Univariate standardization controls for time-variation in the variances. Fig. 6 shows the conditional standard deviations derived from our GJR-GARCH(1,1)- t_r estimation of the DAX30 returns. There is evidence of a long swing




Fig. 5. Exceedance correlations for the CAC40 versus DAX30 returns.




Fig. 6. Conditional standard deviations of the DAX30 returns. Note: The conditional standard deviations are estimated for an asymmetric GJR-GARCH(1,1)- t_ν model with Student- t_ν distributed disturbances.

in volatility with volatility at particularly high levels from 1997 until 2004. The conditional standard deviation figures look very similar for the other index return series, which may well be due to a shift in factor volatility (in a world factor pricing model). See [Bekaert et al. \(2005\)](#). Our GARCH standardization (partially) controls for this phenomenon. However, there may still be an unexplained shift in correlation which needs to be taken into account.

One formal approach to estimate time-varying correlation is to estimate a multivariate GARCH model. Rather than pursuing this computationally intensive strategy, we estimate simple rolling correlations between standardized returns to reveal signs of time-variation in correlation.

[Fig. 7](#) shows rolling correlations (for a rolling window size of 40 daily observations) for the CAC40 versus DAX30 returns. Perhaps surprisingly, correlation seems remarkably stable (albeit at higher levels) during the high volatility period identified in [Fig. 6](#). Similar stability results are obtained by [Chesnay and Jondeau \(2001\)](#), based on a Markov-switching model with regime-dependent correlation. To more formally test this phenomenon, we estimate the dynamic conditional correlations (DCC) following [Engle's \(2002\)](#) methodology. [Fig. 8](#) confirms our initial hunch of a stable correlation sub-period for the CAC40 versus DAX30 returns, based on this DCC measure. We split our sample into two more or less equal sized sub-sample periods to evaluate the impact of time-variation in correlation on our truncated and exceedance correlation results.




Fig. 7. Rolling window correlations. Note: The rolling correlations are estimated over a rolling window of 40 historical observations of daily returns.




Fig. 8. Dynamic conditional correlations for the DAX30 versus CAC40 returns. Note: Dynamic conditional correlations are estimated in two stages, with the first stage consisting of univariate asymmetric GJR-GARCH(1,1)- t models for the DAX30 and the CAC40 returns.

Sample period 1 extends from 2 January 1990 until 20 May 1997, while sample period 2 extends from 21 May 1997 until 3 March 2005. In Table 3, the second period unconditional correlations (between square brackets) are as expected significantly higher than for the first period. In Table 4, we see that the second period Student- t_r degrees of freedom parameter estimates (between square brackets) indicates a significantly fatter tailed bivariate distribution, consistent with the higher levels of volatility.

To extend our earlier conditional correlation results, we computed truncated and exceedance correlations for both sub-sample periods.

We see in Fig. 9 that the assumption of Student- t_r provides a better fit for the downside and upside exceedance correlations for the DAX30 versus CAC40 example in both periods. There is evidence of short run variation in correlations which is only slightly corrected for by the volatility correction in period 1. The period 1 empirical exceedance correlations suggest excessive correlation beyond what can be explained by a bivariate Student- t_r distribution, however in both samples the Student- t_r is more able to capture this variation than the assumption of normality. Although there is substantial time variation in correlation and in correlations conditional on the return magnitude, driven maybe in part from changing factor volatilities, the Student- t_r model is best able to capture the conditional correlation structure for the equity markets analysed.

From a pragmatic point of view the theoretical correlation function, assuming a bivariate Student- t_r with 7 degrees of freedom for the CAC40 and the DAX30 provides an almost flat structure. In both sub-samples, there is little evidence of increasing conditional correlation over the return structure, but definitive evidence of fat-tails. If normality were assumed, then the apparent increase in correlation would simply be due to fat tails in the empirical distribution.


3.4. Goodness of fit tests

Finally, Tables 5 and 6 give the formal asymmetry H statistics for our three conditional correlation estimates against the maintained distribution (respectively bivariate normal, and bivariate Student- t_r in italics).

Table 5 gives the H statistics for the full sample.¹² For the truncated correlations, the H statistics are not significantly different for the two distributional assumptions, reflecting the difficulty in discriminating between the normal and Student- t_r truncated correlations. While the overall H test is rejected for all four index combinations and both assumed distributions, the H^+ , H^- , and AH test are never rejected at the 95% confidence level.

¹² The results are given for the four country index returns against the MSCI World index returns. The results for 'cross'-country index returns are available from the authors.

Exceedance correlations for the DAX30 versus CAC40 returns – Period 1



Exceedance correlations for the DAX30 versus CAC40 returns – Period 2




Fig. 9. Exceedance correlations for the DAX30 versus CAC40 returns — Period 1 and Period 2. Note: Period 1 refers to 2 January 1990 until 20 May 1997. Period 2 refers to 21 May 1997 until 3 March 2005.

However, for the cumulative truncated correlations, we do find significant differences when comparing the distributional assumptions. According to the H^- test, the Student- t , assumption provides a significantly better fit than the normal assumption for the S&P500 versus MSCLWorld combination (the normal assumption is rejected). This outcome is reversed for the H^+ test, suggesting asymmetry in cumulative truncated correlations for bear versus bull market conditions. Accordingly, the overall H test is rejected for both distributional assumptions. For the exceedance

Table 5

H statistics — A joint test on significant excess correlation

	Model	<i>H</i>	<i>H</i> ⁺	<i>H</i> [−]	<i>AH</i>
<i>Truncated correlations</i>					
S&P500	Normal	0.060*	0.046	0.071	0.011
	<i>t</i>	0.085*	0.089	0.080	0.029
FTSE100	Normal	0.052*	0.055	0.048	0.005
	<i>t</i>	0.048*	0.057	0.038	0.015
CAC40	Normal	0.059*	0.055	0.063	0.007
	<i>t</i>	0.063*	0.063	0.062	0.018
DAX30	Normal	0.073*	0.064	0.080	−0.012
	<i>t</i>	0.081*	0.071	0.090	−0.001
<i>Cumulative truncated correlations</i>					
S&P500	Normal	0.046*	0.017	0.063*	−0.025
	<i>t</i>	0.062*	0.083*	0.031	0.035
FTSE100	Normal	0.025	0.013	0.032	−0.014
	<i>t</i>	0.030	0.034	0.026	0.019
CAC40	Normal	0.029	0.025	0.033	−0.005
	<i>t</i>	0.045*	0.059*	0.022	0.030
DAX30	Normal	0.021	0.015	0.025	0.009
	<i>t</i>	0.054*	0.037*	0.067*	0.047*
<i>Exceedance correlations</i>					
S&P500	Normal	0.162*	0.144*	0.178*	−0.153*
	<i>t</i>	0.042	0.022	0.055	−0.036
FTSE100	Normal	0.112*	0.060*	0.145*	−0.094*
	<i>t</i>	0.071*	0.028	0.095*	−0.043
CAC40	Normal	0.130*	0.117*	0.141*	−0.105*
	<i>t</i>	0.085*	0.083*	0.085*	−0.052*
DAX30	Normal	0.090*	0.116*	0.053*	−0.056*
	<i>t</i>	0.064	0.080*	0.043	0.005

The table gives the asymmetry *H* statistics assuming the null hypothesis of a bivariate normal distribution, respectively a bivariate Student-*t*_ν distribution. To compute the *H* statistics, weights are taken proportional to the number of observations in each correlation sample (truncated, cumulative truncated, and exceedance). All correlations are computed for the individual countries' standardized returns against the MSCI World index standardized returns. Asterisks indicate significant rejection of the assumed distribution model at the 95% confidence level.

correlations, the results do not indicate this 'bear-versus-bull' asymmetry. The *H*⁺ and *H*[−] tests are now in agreement. While the normal assumption is rejected, the Student-*t*_ν assumption cannot be rejected for the S&P500, and the DAX30, versus the MSCI World combination.

Overall, from the 48 comparisons of *H* statistics reported in Table 5, 23 return smaller *H* statistics for the Student-*t*_ν than for the normal distribution. However, from the 16 comparisons of *H* statistics for exceedance correlations only, all 16 return smaller *H* statistics for the Student-*t*_ν than for the normal distribution. And, from the 12 comparisons of *H*[−] statistics, 9 return smaller *H*[−] statistics for the Student-*t*_ν than for the normal distribution.

More specifically, Table 6 gives the *H* tests for the two sub-sample periods for the DAX30 versus CAC40 returns. The *H* tests are always rejected for the normal distribution assumption in both periods. The *H* (Student-*t*_ν) tests cannot be rejected for the period 1 cumulative truncated correlations and period 2 exceedance correlations. The sub-*H* tests are mostly rejected in period 1 for both distributional assumptions. The main result, however, is the rejection of the normal assumption for period 2 exceedance correlations, but the apparent good fit of the Student-*t*_ν assumption for upside/downside and overall exceedance correlations. From the 24 comparisons of *H* statistics reported in Table 6, 15 return smaller *H* statistics for the Student-*t*_ν than for the normal distribution. From the 8 comparisons of *H* statistics for exceedance correlations, all 8 return smaller *H* statistics for the Student-*t*_ν than for the normal distribution. And, from the 6 comparisons of *H*[−] statistics, 4 return smaller *H*[−] statistics for the Student-*t*_ν than for the normal distribution. Overall, downside exceedance correlations do indeed seem to be better fit by a Student-*t*_ν distribution than by a normal distribution, after controlling for substantial time variation in volatility and correlation.

4. Conclusions

Turbulent financial market conditions easily lead to an impression of contagion and spillover effects wreaking havoc on the benefits of international diversification. It is intuitively appealing to conclude that correlation between international financial asset returns increases during volatile market conditions. If this intuition is corroborated by empirical evidence, it could indeed have serious implications for international portfolio allocation.

A variety of conditional correlation estimators have recently been proposed to estimate and quantify this effect of increasing correlation, conditional on the size of returns. Unfortunately, most of these estimators imply a conditional correlation structure that cannot be easily compared to the unconditional correlation estimate.

In this paper, we evaluate the performance of two popular conditional correlation estimators, the truncated correlation estimator and the exceedance correlation estimator. We derive theoretical conditional correlation functions for both estimators assuming bivariate normal return distributions. Since the joint conditional normality assumption is, in fact, an unlikely candidate for financial asset returns, we analytically derive both these estimators for the fatter-tailed bivariate Student- t_r distribution. The theoretical conditional correlation functions are distinctly different under the Student- t_r assumption. In fact, these respective theoretical functions suggest that earlier studies may have overestimated the excess in conditional correlation by assuming bivariate normality.

We apply the truncated and exceedance correlation estimators to a data set of international stock market index returns. We investigate the univariate properties of the index returns and then filter them to obtain standardized returns. An asymmetric GJR-GARCH(1,1) specification with Student- t_r distributed disturbances fitted the series best. Under the assumption of normally distributed standardized returns, we find evidence of significant excess conditional correlation in the tails of the bivariate distribution. This would indicate that assuming normality in the (conditional) variance–covariance matrix overestimates the diversification benefits in dynamic mean-variance portfolio allocation. When assuming the Student- t_r distribution, we find that excess conditional correlation frequently disappears, at least for the left tail of the bivariate return distribution. The Student- t_r is better at modelling the conditional correlation structure empirically observed in the data. Our results highlight the fact that correlation breakdown may be a spurious artefact of the assumption of normally distributed returns. We also find significant evidence of asymmetry in the conditional correlation functions. The shape of this asymmetry seems to point towards a mixture of two distributions, potentially driven by a regime-switching model (Ang and Bekaert, 2002) where factor volatility appears to change in different regimes.

Table 6
H Statistics for the sub-sample periods (CAC40 versus DAX30)

	Model	H	H^+	H^-	AH
<i>Period 1</i>					
Truncated correlations	Normal	0.106*	0.111*	0.100*	0.008
	t	0.086*	0.094*	0.078	0.018
Truncated correlations (cumulative)	Normal	0.084*	0.056	0.104*	−0.075*
	t	0.048	0.019	0.065*	−0.036
Exceedance correlations	Normal	0.274*	0.277*	0.270*	−0.265*
	t	0.168*	0.172*	0.164*	−0.162*
<i>Period 2</i>					
Truncated correlations	Normal	0.097*	0.110*	0.085	−0.001
	t	0.135*	0.152*	0.115*	0.017
Truncated correlations (cumulative)	Normal	0.091*	0.124*	0.034	0.067
	t	0.151*	0.191*	0.096*	0.127*
Exceedance correlations	Normal	0.160*	0.134*	0.184*	−0.153*
	t	0.054	0.070	0.031	0.039

The table gives the asymmetry H statistics assuming the null hypothesis of a bivariate normal distribution, respectively a bivariate Student- t_r distribution. To compute the H statistics, weights are taken proportional to the number of observations in each correlation sample (truncated, cumulative truncated, and exceedance). Conditional correlations are computed for the CAC40 versus DAX30 returns. Period 1 refers to 2 January 1990 until 16 May 1997. Period 2 refers to 21 May 1997 until 3 March 2005. Truncated correlation estimates are conditioned on a threshold return for the DAX30. Asterisks indicate significant rejection of the assumed distribution model at the 95% confidence level.

Appendix A. Mean and variance of the truncated Student- t distribution

Let the random variable t have a Student- t distribution with r degrees of freedom and probability density function

$$f_r(t) = C_r \left(1 + \frac{t^2}{r}\right)^{-(r+1)/2}, \quad \text{for } -\infty < t < \infty \quad (\text{A1})$$

with $C_r = \frac{\Gamma[(r+1)/2]}{\sqrt{\pi r} \Gamma(r/2)}$. The distribution of t has mean zero and variance $r/(r-2)$, and cumulative probability distribution function given by

$$F_r(x) = \Pr[t \leq x] = \int_{-\infty}^x f_r(t) dt, \quad \text{for } -\infty < t < \infty. \quad (\text{A2})$$

A.1. Standardized Student- t distribution

Let $\varepsilon = \frac{\sqrt{r-2}}{\sqrt{r}}t$, so that $E[\varepsilon] = 0$ and $\text{var}[\varepsilon] = 1$. We say that ε has a Standardized Student- t distribution with r degrees of freedom (denoted by SS- t_r) and has probability density function given by

$$g_r(\varepsilon) = C_r^{\varepsilon} \left(1 + \frac{\varepsilon^2}{r-2}\right)^{-(r+1)/2}, \quad -\infty < \varepsilon < \infty$$

with

$$C_r^{\varepsilon} = C_r \frac{\sqrt{r}}{\sqrt{r-2}} = \frac{\Gamma[(r+1)/2]}{\sqrt{\pi(r-2)} \Gamma(r/2)}.$$

The cumulative probability distribution function associated with ε is easily derived from F_r as follows:

$$G_r(a) = \Pr[\varepsilon \leq a] = \Pr\left[t \leq \frac{\sqrt{r}}{\sqrt{r-2}}a\right] = F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}a\right), \quad \text{for } -\infty < a < \infty.$$

A.2. $E[\varepsilon | L < \varepsilon \leq U]$

Suppose that interest lies in the moments of the distribution of ε subject to the constraint that $L < \varepsilon \leq U$.

The conditional mean of ε subject to the constraint that $L < \varepsilon \leq U$ is defined via integration of ε with respect to the conditional distribution of ε given $L < \varepsilon \leq U$:

$$E[\varepsilon | L < \varepsilon \leq U] = \int_L^U \frac{\varepsilon g_r(\varepsilon)}{\Pr(L < \varepsilon \leq U)} d\varepsilon.$$

Since

$$\begin{aligned}
 \int_L^U \varepsilon g_r(\varepsilon) d\varepsilon &= C_r^e \int_L^U \varepsilon \left(1 + \frac{\varepsilon^2}{(r-2)}\right)^{-(r+1)/2} d\varepsilon \\
 &= C_r^e \frac{(r-2)}{2} \int_L^U \frac{2\varepsilon}{(r-2)} \left(1 + \frac{\varepsilon^2}{(r-2)}\right)^{-(r+1)/2} d\varepsilon \\
 &= C_r^e \frac{(r-2)}{(r-1)} \left[-\left(1 + \frac{U^2}{(r-2)}\right)^{-(r-1)/2} + \left(1 + \frac{L^2}{(r-2)}\right)^{-(r-1)/2} \right] \\
 &= \frac{C_r^e (r-2)}{C_{r-2} (r-1)} [-f_{r-2}(U) + f_{r-2}(L)] = f_{r-2}(L) - f_{r-2}(U)
 \end{aligned} \tag{A3}$$

and hence

$$E[\varepsilon | L < \varepsilon \leq U] = \frac{[f_{r-2}(L) - f_{r-2}(U)]}{\left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L\right)\right]}$$

A.3. $\text{var}[\varepsilon | L < \varepsilon \leq U]$

To calculate the conditional variance, first consider the second non-central conditional moment:

$$E[\varepsilon^2 | L < \varepsilon \leq U] = \frac{1}{\Pr[L < \varepsilon \leq U]} \int_L^U \varepsilon^2 g_r(\varepsilon) d\varepsilon \tag{A4}$$

The integral in Eq. (A4) may be solved using integration by parts, with

$$\begin{aligned}
 u &= \varepsilon & dv &= \varepsilon g_r(\varepsilon) d\varepsilon \\
 du &= d\varepsilon & v &= -f_{r-2}(\varepsilon),
 \end{aligned}$$

and hence

$$\int_L^U \varepsilon^2 g_r(\varepsilon) d\varepsilon = -\varepsilon f_{r-2}(\varepsilon) \Big|_L^U + \int_L^U f_{r-2}(\varepsilon) d\varepsilon = Lf_{r-2}(L) - Uf_{r-2}(U) + F_{r-2}(U) - F_{r-2}(L) \tag{A5}$$

resulting in

$$E[\varepsilon^2 | L < \varepsilon \leq U] = \frac{Lf_{r-2}(L) - Uf_{r-2}(U) + F_{r-2}(U) - F_{r-2}(L)}{\left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L\right)\right]} \tag{A6}$$

The conditional variance may be obtained from Eqs. (A3) and (A6) using the usual relationship

$$\text{var}[\varepsilon^2 | L < \varepsilon \leq U] = E[\varepsilon^2 | L < \varepsilon \leq U] - \{E[\varepsilon | L < \varepsilon \leq U]\}^2 \tag{A7}$$

Appendix B. Exceedance correlations for the Student- t distribution

Now suppose ε_x and ε_y are two independent SS- t_r random variables and let

$$x = \varepsilon_x \text{ and } y = \rho \varepsilon_x + \sqrt{(1 - \rho^2)} \varepsilon_y$$

Clearly, $E[x]=E[y]=0$, $\text{var}(x)=\text{var}(y)=1$ and the unconditional correlation is $\text{corr}(x,y)=\rho$.

However, we are interested in the correlation between (x,y) when pairs are constrained to lie in the region defined by $Z=\{(x,y)|L<x\leq U, L<y\leq U\}$. That is, we are interested in determining

$$\text{corr}(x,y|Z) = \frac{\text{cov}(x,y|Z)}{\sqrt{\text{var}(x|Z)\text{var}(y|Z)}} = \frac{E[xy|Z] - E[x|Z]E[y|Z]}{\sqrt{\text{var}(x|Z)\text{var}(y|Z)}}. \quad (\text{A8})$$

Note that $(x,y)\in Z$ if and only if $(\varepsilon_x, \varepsilon_y)\in Z^e$, where

$$Z^e = \{(\varepsilon_x, \varepsilon_y)|L<\varepsilon_x\leq U, L^*(\varepsilon_x)<\varepsilon_y\leq U^*(\varepsilon_x)\}$$

and where

$$U^*(\varepsilon) = \frac{(U - \rho\varepsilon)}{\sqrt{(1 - \rho^2)}}, \text{ and } L^*(\varepsilon) = \frac{(L - \rho\varepsilon)}{\sqrt{(1 - \rho^2)}}. \quad (\text{A9})$$

The probability of the event Z is obtained from

$$\Pr(Z) = \int_L^U \int_{L^*(\varepsilon_x)}^{U^*(\varepsilon_x)} g_r(\varepsilon_x)g_r(\varepsilon_y)d\varepsilon_y d\varepsilon_x = \int_L^U g_r(\varepsilon_x) \left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U^*(\varepsilon_x)\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L^*(\varepsilon_x)\right) \right] d\varepsilon_x$$

since $\Pr(Z)=\Pr(Z^e)$. Although the above integral does not appear to be available in closed form, it may be written as an expectation with respect to the (unconstrained) SS- t_r distribution:

$$\Pr(Z) = E\left[1_{(L,U]}(\varepsilon) \left[F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}U^*(\varepsilon)\right) - F_r\left(\frac{\sqrt{r}}{\sqrt{r-2}}L^*(\varepsilon)\right) \right] \right], \quad (\text{A10})$$

where $1_{(L,U]}(\varepsilon)$ denotes the indicator function for the interval $L<\varepsilon\leq U$, defined as

$$1_{(L,U]}(\varepsilon) = \begin{cases} 1 & \text{if } L<\varepsilon\leq U \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A11})$$

From Eq. (A8), it can be seen that the conditional covariance between x and y can be constructed once $E[x|Z]$, $E[y|Z]$, $\text{var}(x|Z)$, $\text{var}(y|Z)$ and $E[xy|Z]$ are determined. The conditional moments associated with x , namely $E[x|Z]$ and $\text{var}(x|Z)$, may be determined directly from the conditional moments of ε_x , as in Eqs. (A3) and (A7). The remaining conditional moments require more effort before they can be determined.

B.1. $E[y|Z]$

Consider first the conditional mean of y given Z

$$E[y|Z] = E[\rho\varepsilon_x + \sqrt{(1 - \rho^2)}\varepsilon_y|Z] = \rho E[\varepsilon_x|Z] + \sqrt{(1 - \rho^2)}E[\varepsilon_y|Z] = \rho E[\varepsilon_x|L<\varepsilon_x\leq U] + \sqrt{(1 - \rho^2)}E[\varepsilon_y|Z]$$

The first term is available using Eq. (A3), but the second term is more complex due to the conditioning on Z . Specifically

$$\begin{aligned} E[\varepsilon_y|Z] &= [\Pr(Z)]^{-1} \int_L^U \int_{L^*(\varepsilon_x)}^{U^*(\varepsilon_x)} \varepsilon_y g_r(\varepsilon_x) g_r(\varepsilon_y) d\varepsilon_y d\varepsilon_x \\ &= \frac{C_r^e(r-2)}{2[\Pr(z)]} \int_L^U g_r(\varepsilon_x) \int_{U^*(\varepsilon_x)}^{L^*(\varepsilon_x)} \frac{2}{(r-2)} \varepsilon_y \left(1 + \frac{\varepsilon_y^2}{(r-2)}\right)^{-(r+1)/2} d\varepsilon_y d\varepsilon_x \\ &= [\Pr(Z)]^{-1} \int_L^U g_r(\varepsilon_x) \left[f_{r-2} \left(\frac{L - \rho \varepsilon_x}{\sqrt{1 - \rho^2}} \right) - f_{r-2} \left(\frac{U - \rho \varepsilon_x}{\sqrt{1 - \rho^2}} \right) \right] d\varepsilon_x \end{aligned}$$

and hence

$$E[\varepsilon_y|Z] = [\Pr(Z)]^{-1} E[1_{(L,U]}(\varepsilon) [f_{r-2}(L^*(\varepsilon)) - f_{r-2}(U^*(\varepsilon))]] \quad (\text{A12})$$

where $\varepsilon \sim \text{SS-}t_r$.

B.2. $E[y^2|Z]$

Similarly, the second non-central moment of y conditional upon Z may be determined as follows:

$$E[y^2|Z] = E[(\rho \varepsilon_x + \sqrt{(1 - \rho^2)} \varepsilon_y)^2|Z] = \rho^2 E[\varepsilon_x^2|Z] + 2\rho \sqrt{(1 - \rho^2)} E[\varepsilon_x \varepsilon_y|Z] + (1 - \rho^2) E[\varepsilon_y^2|Z]$$

Although $E[\varepsilon_x^2|Z]$ is available from Eq. (A6), both $E[\varepsilon_x \varepsilon_y|Z]$ and $E[\varepsilon_y^2|Z]$ need to be expressed in a form more amenable to computation.

Consider first $E[\varepsilon_y^2|Z]$:

$$E[\varepsilon_y^2|Z] = [\Pr(Z)]^{-1} \int_L^U g_r(\varepsilon_x) \int_{L^*(\varepsilon_x)}^{U^*(\varepsilon_x)} \varepsilon_y^2 g_r(\varepsilon_y) d\varepsilon_y d\varepsilon_x,$$

and note that from Eq. (A5) we have

$$\int_{L^*(\varepsilon_x)}^{U^*(\varepsilon_x)} \varepsilon_y^2 g_r(\varepsilon_y) d\varepsilon_y = L^*(\varepsilon) f_{r-2}(L^*(\varepsilon)) - U^*(\varepsilon) f_{r-2}(U^*(\varepsilon)) + F_{r-2}(U^*(\varepsilon)) - F_{r-2}(L^*(\varepsilon))$$

and hence

$$\begin{aligned} E[\varepsilon_y^2|Z] &= [\Pr(Z)]^{-1} E[1_{(L,U]}(\varepsilon) [L^*(\varepsilon) f_{r-2}(L^*(\varepsilon)) - U^*(\varepsilon) f_{r-2}(U^*(\varepsilon)) \\ &\quad + F_{r-2}(U^*(\varepsilon)) - F_{r-2}(L^*(\varepsilon))]] \end{aligned} \quad (\text{A13})$$

Next we consider $E[\varepsilon_x \varepsilon_y|Z]$. We have

$$\begin{aligned} E[\varepsilon_x \varepsilon_y|Z] &= [\Pr(Z)]^{-1} \int_L^U \varepsilon_x g_r(\varepsilon_x) \int_{L^*(\varepsilon_x)}^{U^*(\varepsilon_x)} \varepsilon_y g_r(\varepsilon_y) d\varepsilon_y d\varepsilon_x \\ &= [\Pr(Z)]^{-1} \int_L^U \varepsilon_x g_r(\varepsilon_x) [f_{r-2}(L^*(\varepsilon_x)) - f_{r-2}(U^*(\varepsilon_x))] d\varepsilon_x \end{aligned}$$

and hence

$$E[\varepsilon_x \varepsilon_y | Z] = [\Pr(Z)]^{-1} E[1_{(L,U]}(\varepsilon) [\varepsilon f_{r-2}(L^*(\varepsilon)) - \varepsilon f_{r-2}(U^*(\varepsilon))]] \quad (\text{A14})$$

B.3. $E[xy|Z]$

The only remaining component of Eq. (A8) is

$$E[xy|Z] = E[\varepsilon_x(\rho \varepsilon_x + \sqrt{(1-\rho^2)}\varepsilon_y)|Z] = \rho E[\varepsilon_x^2|Z] + \sqrt{(1-\rho^2)} E[\varepsilon_x \varepsilon_y|Z]$$

The components of this expectation are available from Eqs. (A4) and (A14).

B.4. Computing $\text{corr}(x,y|Z)$

The only remaining issue for computing the conditional correlation, $\text{corr}(x,y|Z)$, is to evaluate the expectations in Eqs. (A10), (A12), (A13) and (A14). However, as all of the expectations are with respect to the same $SS-t_r$ distribution, they may be easily approximated using a Monte Carlo integration approach.

To approximate the required expectations, produce $\{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n\}$, a (pseudo-) random sample from the $SS-t_r$ distribution. Then, calculate

$$\begin{aligned} h_1(\varepsilon_i) &= 1_{(L,U]}(\varepsilon_i) \left[F_r \left(\frac{\sqrt{r}}{\sqrt{r}-2} U^*(\varepsilon_i) \right) - F_r \left(\frac{\sqrt{r}}{\sqrt{r}-2} L^*(\varepsilon_i) \right) \right] \\ h_2(\varepsilon_i) &= 1_{(L,U]}(\varepsilon_i) [f_{r-2}(L^*(\varepsilon_i)) - f_{r-2}(U^*(\varepsilon_i))] \\ h_3(\varepsilon_i) &= 1_{(L,U]}(\varepsilon_i) [L^*(\varepsilon_i) f_{r-2}(L^*(\varepsilon_i)) - U^*(\varepsilon_i) f_{r-2}(U^*(\varepsilon_i)) + F_{r-2}(U^*(\varepsilon_i)) - F_{r-2}(L^*(\varepsilon_i))] \\ h_4(\varepsilon_i) &= 1_{(L,U]}(\varepsilon_i) [\varepsilon_i f_{r-2}(L^*(\varepsilon_i)) - \varepsilon_i f_{r-2}(U^*(\varepsilon_i))] \end{aligned}$$

for each i , using the definitions in Eqs. (A1), (A2), (A9) and (A11). The sum of the sample of the functions h_1 , h_2 , h_3 , and h_4 may then be used to approximate the corresponding expectations as follows:

$$\begin{aligned} |\Pr(Z) = \frac{1}{n} \sum_{i=1}^n h_i(\varepsilon_i), E[\varepsilon_y|Z^e] = \left[\sum_{i=1}^n h_1(\varepsilon_i) \right]^{-1} \sum_{i=1}^n h_2(\varepsilon_i), E[\varepsilon_y^2|Z^e] = \left[\sum_{i=1}^n h_1(\varepsilon_i) \right]^{-1} h_3(\varepsilon_i), \text{ and } E[\varepsilon_x \varepsilon_y|Z^e] \\ = \left[\sum_{i=1}^n h_1(\varepsilon_i) \right]^{-1} \sum_{i=1}^n h_4(\varepsilon_i). \end{aligned}$$

Appendix C. Simulated standard errors

The conditional correlation estimators require independent and identically distributed univariate returns. We first estimate the parameters of a GJR-GARCH(1,1)- t_r asymmetric time-varying variance model with standardized Student- t_r distributed errors for each of two index return time series (R_x and R_y) with sample size T .

$$\begin{aligned} \sigma_{R_{x,t}}^2 &= \omega_x + \alpha_x R_{x,t-1}^2 + \beta_x \sigma_{R_{x,t-1}}^2 + \gamma_x R_{x,t-1}^2 1_{R_{x,t-1} < 0} & R_{x,t} &= \varepsilon_{R_{x,t}} \sqrt{\sigma_{R_{x,t}}^2} & \varepsilon_{R_{x,t}} &\sim SS t_{r_x} \\ \sigma_{R_{y,t}}^2 &= \omega_y + \alpha_y R_{y,t-1}^2 + \beta_y \sigma_{R_{y,t-1}}^2 + \gamma_y R_{y,t-1}^2 1_{R_{y,t-1} < 0} & R_{y,t} &= \varepsilon_{R_{y,t}} \sqrt{\sigma_{R_{y,t}}^2} & \varepsilon_{R_{y,t}} &\sim SS t_{r_y} \end{aligned}$$

This gives $\overline{\omega_{xy}}, \overline{\alpha_{xy}}, \overline{\beta_{xy}}, \overline{\gamma_{xy}}, \overline{r_{xy}}$ parameter estimates.

Second, we estimate the bivariate Student- t_r degrees of freedom parameter $\bar{r}$ and the unconditional ($\bar{\rho}$) and conditional ($\bar{\rho}_Q$) empirical correlations between the two standardized index return series.

$$\bar{\rho} = \overline{\text{corr}}(\overline{\varepsilon_{R_{x,t}}}, \overline{\varepsilon_{R_{y,t}}}) \quad \bar{\rho}_Q = \overline{\text{corr}}(\overline{\varepsilon_{R_{x,t}}}, \overline{\varepsilon_{R_{y,t}}}|Q)$$

Third, at the estimated empirical parameters, we simulate bivariate correlated time series x and y with the same (empirical) sample size T . We generate i.i.d. innovations $\varepsilon_x, \varepsilon_y$ from a bivariate Student- t_r distribution with $\bar{r}$ degrees of freedom and let

$$\begin{aligned} x_t &= \varepsilon_{x,t} \sqrt{\sigma_{x,t}^2} & y_t &= \bar{\rho}(\varepsilon_{x,t} \sqrt{\sigma_{x,t}^2}) + \sqrt{(1 - \bar{\rho}^2)}(\varepsilon_{y,t} \sqrt{\sigma_{y,t}^2}) \\ \sigma_{x,t}^2 &= \bar{\omega}_x + \bar{\alpha}_x x_{t-1}^2 + \bar{\beta}_x \sigma_{x,t-1}^2 + \bar{\gamma}_x x_{t-1}^2 1_{x_{t-1} < 0} & \sigma_{y,t}^2 &= \bar{\omega}_y + \bar{\alpha}_y y_{t-1}^2 + \bar{\beta}_y \sigma_{y,t-1}^2 + \bar{\gamma}_y y_{t-1}^2 1_{y_{t-1} < 0} \end{aligned}$$

Fourth, for each simulation of (x, y) , we estimate the parameters ($\tilde{\omega}_{x,y}, \tilde{\alpha}_{x,y}, \tilde{\beta}_{x,y}, \tilde{\gamma}_{x,y}, \tilde{r}_{x,y}$) of univariate GJR-GARCH(1,1)- t_r models and we estimate the bivariate Student- t_r degrees of freedom parameter $\tilde{r}$ from the standardized residuals.

Fifth, we estimate the simulation's conditional correlations $\tilde{\rho}_Q(x_t, \tilde{\sigma}_{x,t}^{-1} y_t, \tilde{\sigma}_{y,t}^{-1} | Q)$ as well as the distribution implied conditional correlations $\tilde{\rho}_Q(t_{\tilde{r}})$ at the estimated bivariate Student- t_r degrees of freedom parameter $\tilde{r}$ according to Section 2.1.

Sixth, we compute H statistics based on the difference between the simulation's conditional correlations $\tilde{\rho}_Q(x_t, \tilde{\sigma}_{x,t}^{-1} y_t, \tilde{\sigma}_{y,t}^{-1} | Q)$ and the distribution implied conditional correlations $\tilde{\rho}_Q(t_{\tilde{r}})$.

We iterate this simulation process (steps 3 to 6) 10,000 times to obtain simulated distributions for the conditional correlations and for the H statistics under the maintained distribution null hypothesis. We compute the relevant standard errors at the appropriate quantiles of these distributions.

References

- Ang, A., Bekaert, G., 2002. International asset allocation with regime shifts. *The Review of Financial Studies* 1137–1187 (Fall).
- Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. *Journal of Financial Economics* 63, 443–494.
- Bekaert, G., Hodrick, R.J., Zhang, X., 2005. International stock return comovements. NBER Working Paper, vol. W11906.
- Barnett, V., 1976. The ordering of multivariate data. *Journal of the Royal Statistical Society. Series A* 139, 318–354.
- Blyth, S., 1996. Out of line. *RISK Magazine* 10 (11), 82–84.
- Bollerslev, T., 1987. A conditionally heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* 69, 542–547.
- Bollerslev, T., Engle, R.F., Wooldridge, J., 1988. A capital asset pricing model with time varying covariances. *Journal of Political Economy* 96, 116–131.
- Boyer, B.H., Gibson, M.S., Loretan, M., 1999. Pitfalls in tests for changes in correlations. *International Finance Discussion Papers*, vol. 597R. Board of Governors of the Federal Reserve Board, Washington, D.C.
- Butler, K.C., Joquin, D.C., 2002. Are the gains from international portfolio diversification exaggerated? The influence of downside risk in bear markets. *Journal of International Money and Finance* 21 (7), 981–1011.
- Campbell, R., Koedijk, K., Kofman, P., 2002. Increased correlation in bear markets. *Financial Analysts Journal* 58 (1), 87–94.
- Chesnay, F., Jondeau, E., 2001. Does correlation between stock returns really increase during turbulent periods? *Economic Notes* 30 (1), 53–80.
- Doornik, J.A., Hansen, H., 1994. A practical test for univariate and multivariate normality. Discussion Paper. Nuffield College.
- Embrechts, P., McNeil, A., Straumann, D., 1999. Correlation: pitfalls and alternatives. *RISK Magazine* 12 (5), 69–71.
- Engle, R.F., 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20, 339–350.
- Forbes, K.J., Rigobon, R., 2002. No contagion, only interdependence: measuring stock market comovements. *Journal of Finance* 57, 2223–2261.
- Glosten, L.R., Jagannathan, D., Runkle, R., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.
- Granger, C.W.J., Teräsvirta, T., Patton, A.J., 2006. Common factors in conditional distributions for bivariate time series. *Journal of Econometrics* 132, 43–57.
- Harvey, C.R., Siddique, A., 2000. Conditional skewness in asset pricing tests. *Journal of Finance* 55, 1263–1295.
- Huisman, R., Koedijk, K.G., Pownall, R.A., 1998. VaR-x: Fat tails in financial risk management. *Journal of Risk* 1, 43–61.
- Jansen, D.W., De Vries, C.G., 1991. On the frequency of large stock returns: putting booms and busts into perspective. *Review of Economics and Statistics* 73, 19–24.
- Jondeau, E., Rockinger, M., 2003. Testing for differences in the tails of stock-market returns. *Journal of Empirical Finance* 10, 559–581.
- Johnson, N.L., Kotz, S., 1972. *Distributions in Statistics: Continuous Multivariate Distributions*. John Wiley & Sons, Inc.

- Karolyi, G.A., 1995. A multivariate GARCH model of international transmissions of stock returns and volatility: the case of the United States and Canada. *Journal of Business and Economic Statistics* 13, 11–25.
- Karolyi, G.A., Stulz, R.M., 1996. Why do markets move together? An investigation of U.S.–Japan stock return comovements. *Journal of Finance* 51 (3), 951–986.
- Longin, F., 2005. The choice of the distribution of asset returns: how extreme value theory can help? *Journal of Banking and Finance* 29, 1017–1035.
- Longin, F., Solnik, B., 1995. Is the correlation in international equity returns constant: 1960–1990? *Journal of International Money and Finance* 14, 3–26.
- Longin, F., Solnik, B., 2001. Extreme correlation of international equity markets. *Journal of Finance* 56 (2), 649–676.
- Loretan, M., English, W.B., 2000a. Evaluating “correlation breakdowns” during periods of market volatility. *International Finance Discussion Papers*, vol. 658. Board of Governors of the Federal Reserve Board, Washington, D.C.
- Loretan, M., English, W.B., 2000b. III. Special feature: evaluating changes in correlations during periods of high market volatility. *BIS Quarterly Review* 29–36 (June).
- Newey, W.K., West, K.D., 1987. A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Patton, A.J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *Journal of Financial Econometrics* 2, 130–168.
- Sullivan, G., 1995. Correlation counts. *RISK Magazine* 8, 36–38.