



**DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS**

**Diversification meltdown or the impact of fat tails on
conditional correlation?**

**Rachel Campbell, Catherine S. Forbes,
Kees Koedijk and Paul Kofman**

Diversification meltdown or the impact of fat tails on conditional correlation?*

Rachel Campbell
Maastricht University

Catherine S. Forbes
Monash University

Kees Koedijk,
Erasmus University Rotterdam, CEPR

Paul Kofman
The University of Melbourne

27 November 2003

Abstract

A perceived increase in correlation during turbulent market conditions implies a reduction in the benefits arising from portfolio diversification. Unfortunately, it is exactly then that these benefits are most needed. To determine whether diversification truly breaks down, we investigate the robustness of a popular conditional correlation estimator against alternative distributional assumptions. Analytical results show that the apparent meltdown in diversification could be a result of assuming normally distributed returns. A more realistic assumption – the bivariate Student- t distribution – suggests that there is little empirical support for diversification meltdown.

JEL classification: G11, G14

Keywords: Conditional correlation; Truncated correlation; Bivariate Student- t correlation;

* All errors pertain to the authors. 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 and Steve Satchell and Frans de Roon for valuable comments. The Australian Research Council is gratefully acknowledged for financial support.
Corresponding author: R. Campbell, Faculty of Economics and Business Administration, Limburg Institute of Financial Economics (LIFE), Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands, Tel. +31-43-388 3648 / 3838, Fax. +31-43-388 4875, Email: r.campbell@berfin.unimaas.nl

1. Introduction

Investors' ability to accurately estimate (and forecast) the correlation between financial asset returns is the key to successful portfolio management. Indeed, the benefits arising from diversifying risk make it worthwhile to invest part of the portfolio in assets which offer – at first sight – inferior expected returns. However, both market lore (e.g., Sullivan, 1995 and Blyth, 1996) and recent academic research (Longin and Solnik, 1995 and Karolyi and Stulz, 1996, Kritzman et al., 2001) suggest that these benefits rapidly erode during turbulent market conditions. Large (extreme) movements in financial asset prices are found to be more highly correlated than moderate/small movements. In particular, large falls in international financial market prices occur with greater simultaneity than the assumption of constant correlation would predict (Ang and Chen, 2002, and Longin and Solnik, 2001). Since it is precisely under these conditions that diversification is needed most, investors should be extremely concerned about “correlation breakdown”¹.

Assume that correlation does indeed depend on the size of asset returns. To avoid confusion in terminology, we should carefully define what we mean with size-dependent (or, size-conditional) correlation. We distinguish predictable size-dependent correlation from random size-dependent correlation. Predictable size-dependency in correlation arises, for example, in multivariate GARCH models. The well-known empirical feature of intertemporal persistence in the volatility, or size, of returns will also imply persistence in their correlation. In effect, time-dependent volatility (correlation) implies size-dependent volatility (correlation). Investors would then be wise to condition their forecast correlations on recent conditional correlations when optimally allocating their portfolios. But, is this true for all investors? More specifically, do the conditional correlation forecasts accurately indicate the ex-post diversification benefits? The answer depends on the investment horizon. An example

¹ A term coined by Boyer, Gibson, and Loretan (1999). Since correlation actually strengthens under this scenario, we prefer to use the term “diversification meltdown”.

may illustrate this. Consider an investor whose portfolio allocation on the 1st of January 2003 is based on a sample period of historical returns characterized as a persistent bear market (e.g., 2000-2002). Since correlation is positively related to the size of returns, our investor will estimate high correlations between individual asset returns. Are these good (i.e., unbiased) forecast correlations, and hence relevant for the duration of the investment? With persistence in conditional correlation, short horizon investors can reasonably extrapolate the estimated correlations. But what about intermediate and long horizon investors? Eventually, the persistent bear market will give way to more conventional market conditions. At longer horizons, the persistence in, and hence predictability of, correlation diminishes. The sample of historical returns will no longer be representative of the relevant joint distribution of returns. The investor would then like to know how a (biased) conditional correlation estimated from a select sample can still provide inference regarding the relevant population correlation.

Now imagine that our investor somehow obtains a much longer sample of historical returns (perhaps dating back to 1990 covering bear and bull market episodes). A multivariate GARCH model could now reliably be estimated and the short-horizon investors could feasibly use the predictable conditional correlations. Would the unconditional correlation estimate be a sufficient measure to forecast correlation for our long-horizon investor? Not necessarily, if there is also evidence of random size-dependent correlation, i.e., size-dependent correlation in the standardized returns. To investigate the existence of random size-dependency, we need to introduce a conditional correlation estimator.

A number of conditional correlation estimators have recently been proposed to investigate the intertemporal and random size dependency of correlation between asset returns. We investigate the robustness of one particular size-conditional correlation estimator, the truncated correlation estimator, which has recently been applied by several authors to

investigate the robustness of international portfolio diversification benefits. Butler and Joaquin (2002) find that, after accounting for theoretical bias in conditional correlation estimates, correlation does indeed depend on the size of returns. However, this conclusion is based on the assumption that asset returns are jointly normally distributed. Clearly a questionable assumption given abundant empirical evidence of fat-tailedness in the (standardized) asset return distributions. Since our interest is precisely in ‘tail’ correlations (where the benefits from diversification are most needed), we derive a truncated correlation estimator for (one popular class of) fat-tailed return distributions. We find that earlier results supporting diversification meltdown no longer hold when the underlying returns are jointly Student- t distributed. This suggests that long-horizon portfolio managers need not necessarily worry about size-dependent correlation. That is, as long as their portfolio allocation methodology acknowledges fat-tailed return distributions, the unconditional correlation assumption could still be maintained. We also derive an implied unconditional correlation estimator that allows us to infer the population correlation from truncated correlation estimates due to select samples (characterized as bear or bull market episodes). This implied estimator is easier to interpret and more straightforward to implement in portfolio allocation than its truncated equivalent.

The outline of the paper is as follows. In the following section we briefly discuss the estimation methodology for the truncated correlation estimator when the standardized asset returns are jointly normally distributed. We then derive the analogue for standardized asset returns that are jointly Student- t distributed. This estimator includes the joint normal distribution as a limiting case (when the degrees of freedom approximate infinity). We illustrate the theoretical bias in truncated correlation estimates by computing correlation functions for a range of t -degrees of freedom where the population correlation is independent

of size. We also discuss the computation of appropriate standard errors. In Section 3, we apply the truncated correlation estimator to daily data on various international stock market index returns. We note the importance of first standardizing the returns to remove the intertemporal dependency, and only then computing the empirical truncated correlations. We compare these empirical correlations to their theoretical counterparts (under the null of size-independent correlation) for the distributional assumption that best fits the standardized returns. The implications for portfolio optimisation are discussed in Section 4, and Section 5 concludes.

2. Methodology

The liberalisation of capital flows and integration of financial markets are generally considered to be contributing factors to an increase in the correlation between international financial asset returns. While emerging markets' excess returns become a thing of the past, the benefits of international diversification disappear as well; see Butler and Joaquin (2002)². In addition to this 'evolutionary' increase in correlation, there is also some evidence that correlation (between markets and between asset classes) increases during turbulent market conditions. Contagion between international financial markets causes large price movements in one market to spillover into other international financial markets. It therefore seems reasonable to suspect a link between the long-term integration of international financial markets and the increased likelihood of return spillovers. Many papers – including Engle, Lin and Ito (1994) and Karolyi and Stulz (1996) – have focussed on the issue of stock market spillovers implied by increasing correlation between international asset returns during especially volatile market conditions.

² McDonald (2000) claims that these diversification benefits never existed in the first place!

Despite the fact that many practitioners believe that correlation changes intertemporally, and depends on the level of volatility, empirical evidence for this phenomenon is somewhat mixed. The two dominant empirical approaches in the correlation literature are those that condition correlation on time, and those that condition correlation on the size of the returns. The (multivariate) GARCH model, in turn, dominates the time-conditional correlation literature, see e.g., Bollerslev, Engle and Wooldridge (1988), Engle and Kroner (1995), and Engle (2002)³. Various multivariate GARCH specifications have been proposed, some of which put restrictions on the covariances (and hence correlations) to reduce the inevitable parameter dimensionality problem. These restrictions are clearly not appropriate if the primary interest is in the stochastic behaviour of correlation. The size-conditional literature can be further segmented into the *extreme value theory* (EVT) literature, see e.g., Longin and Solnik (2001); the closely related *copula* literature, see e.g., Embrechts, McNeil and Straumann (1999) and Patton (2001); and the *truncated correlation* literature, see e.g., Loretan and English (2000a), Butler and Joaquin (2002), and Forbes and Rigobon (2002). The time- and size-conditional models are certainly not mutually exclusive, a fact that potentially confuses the reader (and practitioner). Patton (2001), e.g., illustrates how the copula approach can be used to model time-varying conditional distributions. Similarly, by conditioning correlation on time, the multivariate GARCH specification links conditional volatility (including covariances) to past volatility. Large returns (of either sign) are followed by large returns, and vice versa. Of course, this implies that time-dependent volatility is also size-dependent. Note, however, that this type of conditional correlation will vary with, and may persist over time but will ultimately mean-revert to its long-term unconditional level. Hence, short-horizon investors will have to be wary of changes in correlation when volatility increases. Long-horizon investors, on the other hand, can safely ignore the intertemporal

³ A stochastic correlation alternative is given in Ball and Torous (2000).

variability in correlation. For them, the correlation between the standardized return series is of greater relevance. However, there is empirical evidence that even after standardizing the return series for GARCH time-dependency, multivariate fat-tailedness persists. Is it possible that the correlation of these standardized returns still varies with size? Both short- and long-horizon investors would then have to reconsider the appropriateness of mean-variance based portfolio allocation. We revisit the practical consequences of size-conditional correlation in Section 4.

The multiplicity of techniques and approaches to estimate conditional correlation suggests the absence of a unique characterization of conditional correlation. Not surprisingly, there are as many conditional correlation estimators as there are different ways of conditioning on size in multivariate return distributions, see Barnett (1976). Rather unsatisfactory for practitioners, the different estimation methodologies do not typically refer to (nor benchmark against) each other. Ang and Chen (2002) show that most of these conditioning schemes cause a bias in the conditional correlation estimates⁴, though not necessarily in the same direction or of the same magnitude. Consider, for example, that we want to measure the correlation between asset returns x and y during bear market conditions. We classify a bear market return as a return below some threshold value λ . The extreme value approach then conditions on joint marginal thresholds, i.e., $\rho(x, y | x < \lambda, y < \lambda)$. For a bivariate normal distribution, this conditional correlation estimator will tend to zero for jointly decreasing marginal thresholds. In contrast, the truncated correlation approach conditions on a single marginal threshold, i.e., $\rho(x, y | x < \lambda)$. For the same bivariate normal distribution, the conditional correlation will then increase for a decreasing marginal

⁴ Campbell, Koedijk and Kofman (2002) develop a conditional correlation estimator that is invariant against conditioning and hence does not suffer from this bias.

threshold⁵. For portfolio allocation purposes, a possibly more intuitive approach would be to condition on a portfolio threshold, $\rho(x, y | x + y < \lambda)$. Campbell, Koedijk and Kofman (2002) show that this scheme is in fact bias-free.

Obviously, the specific type of conditioning will depend on the research question. From a practical perspective, before embarking on a quest for a robust conditional correlation estimator one needs to carefully define the purpose this estimator is supposed to serve. A particular purpose may imply a conditioning scheme and hence, a conditional correlation estimator. Boyer, Gibson and Loretan (1999) discuss conditioning events (e.g., sub-sampling high volatility months) that imply a single marginal conditioning of the joint return distribution⁶. Their conditional correlation estimator conditions on a ‘slice’ of the joint return distribution over which it is estimated. Two examples (sub-samples A and B) of this approach are given in Figure 1. First, consider the case where the market, with return x , has been in a persistent slump and an investment advisor uses short samples to estimate betas for stock y , where $x \in A$. How should we interpret this estimate of beta? Is the beta estimate suitable for both short-term and long-term investors? Or, alternatively, consider the case where x is the return on a guaranteed minimum return portfolio, and y is the return on a hedge portfolio. How do we measure the relevant correlation between x and y given $x \in B$?

INSERT FIGURE 1

It is well known (see e.g., Ang and Chen, 2002) that truncating the joint return distribution according to A (or B) causes bias in the conditional correlation estimate. We need to ‘unbias’ the truncated correlation estimate for the amount and the location of the truncation so that we can properly compare it to the unconditional correlation measure. It is relatively

⁵ Interestingly, the estimation bias is invariably negative, i.e., conditional correlation is always less than unconditional correlation for normally distributed returns! Nevertheless, the bias may be increasing or decreasing when moving into the tails, depending on the conditioning scheme and/or estimator.

⁶ In Campbell, Koedijk and Kofman (2002) we develop an alternative estimator that conditions on portfolio returns instead of the univariate thresholds in this truncated correlation estimator.

straightforward to derive the bias in the theoretical truncated correlation estimator for a bivariate normal distribution. Unfortunately, financial asset returns do not easily fit the normality assumption. Even after allowing for time-dependent variance, the standardized returns frequently display signs of fat-tailedness in excess of normality. This remaining fat-tailedness is often captured by dropping the normality assumption in favour of a Student- t specification, see e.g., Huisman, Koedijk and Pownall (1998). Butler and Joaquin (2002) numerically illustrate how ‘abnormal’ fat-tailedness affects conditional correlation estimates, but to the best of our knowledge an analytical derivation is not currently available for the Student- t case. We therefore extend the truncated correlation estimator for the fat-tailed bivariate Student- t distribution (with degrees of freedom r).

Section 2.1 briefly revisits the truncated correlation estimator for normally distributed returns. We also provide a useful ‘inversion’ result to operationalize the truncated correlation estimates for practical portfolio allocation purposes. We then develop the analogue of the truncated correlation estimator for Student- t distributed returns in Section 2.2. We illustrate the differences in truncated correlation estimates between these maintained distributional hypotheses, based on a series of simulation experiments in Section 2.3.

2.1. Truncated correlation for the bivariate normal distribution

Choose x, y to be correlated random variables driven by independent standard normally distributed (*SND*) random variables $\varepsilon_x, \varepsilon_y$ with drift rates μ_x, μ_y and standard deviations σ_x, σ_y such that,

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

We are interested in the correlation for a partitioning $Q (x, y \mid L \leq x \leq U)$ of the complete bivariate distribution. We can then write the truncated correlation estimator

$$\rho_Q = \frac{\sigma_{xy|Q}}{\sigma_{x|Q}\sigma_{y|Q}} \quad (2)$$

as a ratio of truncated covariance and truncated standard deviations. After some manipulation (see Johnson and Kotz, 1972 p.112) it follows that

$$\rho_Q = \frac{\rho}{\sqrt{\rho^2 + (1 - \rho^2) \frac{\sigma_x^2}{\sigma_{x|Q}^2}}} \quad (3)$$

where ρ is the unconditional correlation between x and y of the complete bivariate distribution. We label the correlation estimator in (3) as the *truncated correlation* estimator. In (3), the truncated variance of x is equivalent to the variance of a truncated normal distribution. The quantiles defining Q are $[L, U]$ such that the truncated variance of x is given by

$$\sigma_{x|Q}^2 = 1 - \left[\frac{\varphi(L) - \varphi(U)}{\Phi(U) - \Phi(L)} \right]^2 + \left[\frac{L\varphi(L) - U\varphi(U)}{\Phi(U) - \Phi(L)} \right] \quad (4)$$

and assuming that the truncation limits are evaluated under the standard normal *pdf* φ and standard normal *cdf* Φ :

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

where p_Q is the probability mass of Q , i.e., $p_Q = \Phi(U) - \Phi(L)$. Some observations are noteworthy. First, if the unconditional variance of x equals its truncated variance, then the truncated correlation equals the unconditional correlation. Second, by truncating the unconditional distribution of x , the variance ratio $(\sigma_x^2 / \sigma_{x|Q}^2)$ will exceed one, which implies that the truncated correlation will be less than the unconditional correlation⁷. For increasing

⁷ There is an exception to this rule when $Q = \{x, y \mid x \in \langle -\infty, L \rangle \cup [U, \infty \rangle\}$, a union of two partitionings as in Loretan and English (2000b), a “high-volatility” partitioning. In that case the variance ratio will 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 exclude this union of partitionings.

truncation, truncated variance, and hence truncated correlation, will decrease monotonically (but non-linearly).

If we partition the unconditional distribution into equal parts (say, deciles), truncated variance will be smaller for the dense central deciles than for the more dispersed tail deciles. Since truncated correlation decreases monotonically in the variance ratio, we can postulate a U-shaped function of truncated correlation vis-à-vis the empirical domain of the marginal distribution of x . The larger the unconditional correlation, ρ , the less pronounced this U-shape will be.

It follows that we can logically expect truncated correlation to increase if we condition on the tails of an empirical distribution in comparison with a central partition of this distribution. Any such increase is, of course, purely spurious. Of some importance is the fact that for finer and finer partitions Q , the variance ratio will increase and hence reduce truncated correlation well below the unconditional correlation. This causes a downward shift in the U-shaped truncated correlation function.

Recall that the truncated correlation estimate will still be less than the unconditional correlation! Empirical truncated correlation estimates can therefore only be compared in a meaningful way with their theoretically implied truncated correlation in (3). If we want to compare empirical truncated correlation estimates with the unconditional correlation, we first need to invert equation (3)

$$\hat{\rho}^{implied} = \sqrt{\frac{\hat{\rho}_Q^2 \hat{\sigma}_x^2}{\hat{\rho}_Q^2 \hat{\sigma}_x^2 + (1 - \hat{\rho}_Q^2) \hat{\sigma}_{x|Q}^2}} \quad (3a)$$

From the truncated correlation estimate, we derive an implied unconditional correlation that is directly comparable to standard unconditional correlation. These implied unconditional correlation estimates would be better suited for the practical implementation discussed in Section 4.

2.2. Truncated correlation for the bivariate Student- t distribution

By comparing theoretical truncated correlation in (3) with unconditional correlation, we are able to identify the theoretical distortion (bias) in the truncated correlation estimator. We can then compare empirical truncated correlations to the relevant theoretical truncated correlations to decide whether correlation is indeed size-dependent. Of course, to make this a valid exercise, the empirical returns will have to satisfy the assumptions underlying the truncated variance of x in (3). A problem arises in using (5) for this purpose if the data are not normally distributed, but instead have fatter tails than normality implies. This would naturally lead to a further dispersion in the tail observations, and to an even steeper increase in truncated correlation. Hence, it would give the mistaken impression of size-conditional correlation even if correlation were inherently size-independent.

According to Boyer, Gibson and Loretan (1999), as long as the bivariate density is elliptic, truncated correlation is defined as in (3). Hence, for a bivariate Student- t where the marginals have identical degrees of freedom, this will hold. Of course, the truncated variance expression in (5) will change even for elliptic distributions. If we assume that the underlying density is jointly Student- t distributed, then the truncated variance becomes:

$$\sigma_{x|Q,r}^2 = \frac{rC_r}{(r-1)P_{[L,U]}^r} \left[L \left(1 + \frac{L^2}{r}\right)^{-(r-1)/2} - U \left(1 + \frac{U^2}{r}\right)^{-(r-1)/2} + \frac{(\sqrt{r})P_{[L,U^*]}^{r-2}}{(\sqrt{r-2})C_{r-2}} \right] - \left[\frac{rC_r}{(r-1)P_{[L,U]}^r} \left[\left(1 + \frac{L^2}{r}\right)^{-(r-1)/2} - \left(1 + \frac{U^2}{r}\right)^{-(r-1)/2} \right] \right]^2 \quad (5a)$$

with

$$C_r = \frac{\Gamma((r+1)/2)}{(r\pi)^{1/2}\Gamma(r/2)}, \quad L^* = \frac{L\sqrt{r-2}}{\sqrt{r}}, \quad U^* = \frac{U\sqrt{r-2}}{\sqrt{r}}$$

where Γ is the gamma function and r is the degrees of freedom parameter of the Student- t distribution. $P_{[L,U]}^r$ is then defined as the probability mass of the interval Q , bounded by U and L , for a Student- t distribution with r degrees of freedom. The derivation of (5a) is given in an appendix to this paper. The expression in (5a) is still computationally straightforward. In graphical terms it also implies a U-shaped truncated correlation function. The increased tail dispersion of the Student- t (in comparison with the normal) distribution generates a steeper U (than for the normal).

2.3. Truncated correlation functions and simulated standard errors

We can now illustrate the impact of fat-tailedness on the theoretical truncated correlation estimates. Truncated correlation functions are computed for two distributional models: the bivariate normal and the bivariate Student- t with a range ($r = 4, 8, 12$) of degrees of freedom. To emphasise the differences between distributional assumptions, we choose a rather large unconditional correlation, $\rho = 0.75$. We partition the bivariate distribution on the x -domain into 20 equal-sized percentiles (5% each) and compute for each percentile the theoretical truncated correlation based on equation (3) and (5) respectively (5a). To allow for sampling error in our empirical experiments we also compute the 95% confidence interval around the theoretical truncated correlation functions. Using equations (1) and (2), we simulate a bivariate return distribution for a sample of size 2,580 (which is identical to our empirical sample size). To obtain appropriate standard errors, we repeated this simulation exercise 1000 times and computed the 2.5 and 97.5 percent confidence limits for the theoretical truncated correlations.

INSERT FIGURE 2

Figure 2 illustrates the theoretical truncated correlation function for a bivariate normal distribution, and a Student- t distribution with $r = 4$ degrees of freedom, with unconditional

correlation assumed equal to $\rho = 0.75$. The Student- t truncated correlations are, as expected, larger than their normal equivalents for the tail percentiles. Note that for the central percentiles, the Student- t truncated correlations are smaller than their normal equivalents. With increasing (decreasing) unconditional correlation, the functions shift upwards (downwards) and the U-shape becomes more (less) pronounced, see also Butler and Joaquin (2002) who illustrate this for the normal distribution assumption.

A Student- t with $r = 4$ degrees of freedom is excessively fat-tailed, but frequently found to feasibly fit empirical asset returns. For higher degrees of freedom, the Student- t becomes more normal and the distinction between the normal and Student- t theoretical truncated correlation functions diminishes. In fact, the distinction is already difficult for the central percentiles. There is, however, an alternative illustration to emphasize the truncated correlation differences for the two distributional assumptions. Instead of considering non-overlapping percentiles, we can also compute theoretical truncated correlations for cumulative percentiles (or rather, decumulative). That is, we first split the x -domain into halves and then estimate the truncated correlation for the lower ($x,y \mid x < \mu_x$) and the upper ($x,y \mid x > \mu_x$) halves, respectively. We then reduce each partitioning successively, thereby gradually moving into the left, respectively right tail percentiles. In (5) and (5a) this implies that we fix the level of L (U) and successively decrease (increase) the level of U (L). Using a Monte Carlo experiment, we simulate the 95% confidence interval around the theoretical truncated correlation functions based on these “shrinking” percentiles. The results are given in Figure 3.

INSERT FIGURE 3

An interesting trade-off occurs. By ‘decumulating’, we reduce the percentile size (and hence its truncated variance). At the same time the remaining percentile becomes more diffuse, since we move into the tail of the distribution, and hence the truncated variance increases. For

the normal distribution (with exponentially declining tails) the first effect more than offsets the second effect. This generates an inverse U-shape for the truncated correlation function. For the Student- t distribution (with tails declining by a power) the second effect more than offsets the first effect when the Student- t is parameterised with less than (about) 5 degrees of freedom. This generates a U-shape for the truncated correlation function with very heavy-tailed distributions, but inverse U-shaped truncated correlation functions for lighter-tailed Student- t distributions. As the degrees of freedom increase for the Student- t (in the limit the distribution approaches normality) the distribution becomes less diffuse, resulting in a smaller truncated variance and hence lower truncated correlation in the tails. Hence, when increasing the degrees of freedom r from say 4 to 12, we will first observe a U-shaped truncated correlation function, which at some level r becomes inverted like the normal distribution shape.

INSERT FIGURE 4

Figure 4, panel A, illustrates the rather wide 95 percent confidence intervals that surround the theoretical correlation functions for the normal and Student- t ($r = 4$) distributions. Standard errors increase for the tail percentiles and are a decreasing function of the sample size. Figure 4, panel B, illustrates how the simulated standard error of a single 5 percent left tail truncated correlation estimate decreases with increasing sample size. The vertical line indicates the sample size of 2,580 in our empirical application. Note that the effective sample size in our examples is only $0.05*N$ (where N is the total sample size). That means that the truncated correlation estimate estimate of the 5 percent left tail percentile is based on only 129 observations, and therefore has large standard errors.

3. Empirical results

The graphs in Section 2.3 clearly show the bias in the theoretical truncated correlation functions for the bivariate normal and bivariate Student- t distributions. They illustrate how easy it is to mistakenly conclude that correlation increases in the tails, particularly when the underlying distribution is fat-tailed. It is therefore vital that we disentangle the spurious increases from the empirically observed tail correlations before deriving any conclusions regarding size-dependent correlation. The relationship in (3) combined with (5) or (5a) allows us to investigate whether there is empirical evidence for diversification meltdown during turbulent market conditions, after discarding truncation bias.

Our data set consists of daily stock market index data (and one bond market index) collected from Datastream for the USA, UK, France, and Germany. The sample period extends from January 1990 to December 1999, i.e., 2581 daily observations. We note that this sample period covers a variety of bull market episodes, bear market episodes and conventional market episodes. The sample is also comparable to the data set used in Longin and Solnik (2001), but our data is based on a higher sampling frequency⁸. We argue that efficient estimation of tail correlations requires as many observations as possible, cf., the standard errors in Figure 4, lower panel.

Using continuously compounded returns on the S&P500, FTSE100, CAC40, DAX100 and the 10-year US Datastream Government Bond Index, we observe that the average return on equity markets averaged about 13.5% over the sample period, twice the return on US Government Bonds. At the same time, the returns on the equity indices were more than twice as volatile as the US Government Bond returns. Summary statistics for the data are given in Table 1.

⁸ Longin and Solnik (2001) use monthly data.

INSERT TABLE 1

Three of the return series exhibit highly significant excess kurtosis, and all but the FTSE100 exhibit significant negative skewness. It is no surprise therefore, that the Jarque-Bera normality test is strongly rejected for every series, except for the FTSE100. This confirms the typical characteristic of asset return distributions with excessive probability mass in the tails relative to the normal distribution. Also note that this excess probability mass seems asymmetrically distributed. Thus, a greater probability of larger (negative) movements in the stock and bond markets than the assumption of normally distributed returns would predict. As we observed in Section 2.3, deviations from normality may have implications for the truncated correlation function of the bivariate returns distribution. The theoretical results in Section 2.2 require that the observations are independently and identically distributed. In keeping with the literature, it seems reasonable to expect that the variance of the asset returns is in fact time-dependent. Similar to Engle (2002), we therefore first filter the univariate series by estimating univariate GARCH(1,1) models and compute the standardized residuals. The GARCH(1,1) parameters are invariably highly significant (also for the US Government Bond returns)⁹. Nevertheless, the GARCH models only capture a limited amount of the observed fat-tailedness in the raw returns. For all standardized return series, the Jarque-Bera test for normality is still rejected.

To gain additional insight into the tail characteristics of the univariate standardized return distributions, we estimate the degrees of freedom parameter r to parameterise the Student- t distribution for each series individually. Maximum likelihood estimation indicates that the series with the higher degree of excess kurtosis tend to have lower estimates for the Student- t degrees of freedom parameter r , with estimates ranging from $\hat{r} = 3.5$ to $\hat{r} = 5$. Each standardized series (including the FTSE100) therefore has considerably fatter tails than the

normal distribution, for which the degrees of freedom tend to infinity. We also estimate the so-called tail index parameter α (see Hill, 1975). For a normal distribution, the tails decline exponentially and the tail index parameter tends to infinity. For fat-tailed distributions, the tails decline by a power and the tail index parameter goes to zero. For the Student- t distribution, the tail index parameter α has the attractive property that it equals the degrees of freedom parameter r . We observe that the tail index parameter estimates are, indeed, satisfactorily close to the Student- t degrees of freedom parameter r estimates.

These empirical findings, confronted with the theoretical results in Section 2, suggest that we are more likely to find prima facie (spurious) evidence of diversification meltdown. Before estimating truncated correlations, $\hat{\rho}_Q$, we need to estimate their unconditional correlation, $\hat{\rho}$. These unconditional correlation estimates are given for the standardized (and raw) returns in Table 2. The unconditional correlation estimate between the standardized S&P500 and the three standardized European return series averages 0.31. The estimated unconditional correlation is much higher between the standardized European return series individually, averaging 0.59. This greater co-movement between standardized European stock market returns implies that 35% (0.59 squared) of stock price movements are common to European markets, whereas 10% of stock price movements are common to both the US and European markets. Not surprisingly, unconditional correlation is less between stock market returns and bond market returns.

INSERT TABLE 2

If we assume that the bivariate distributions are normal, then these unconditional correlation estimates, and the (non-reported) variance ratio estimates, are sufficient to compute the theoretical truncated correlations in equation (3). If the bivariate distribution is better fit by a

⁹ We also estimate EGARCH(1,1) models to allow for asymmetry, and use these standardized residuals whenever the EGARCH(1,1) model outperforms the GARCH(1,1) model.

Student- t , then we first need to estimate the joint degrees of freedom parameter. These parameter estimates are given in Table 3. The range of the joint degrees of freedom parameter estimates ($\hat{r} = 3.5$ to $\hat{r} = 6.3$) is slightly greater than the range of the univariate degrees of freedom parameter estimates, in Table 1.

INSERT TABLE 3

The empirical truncated correlations are estimated for the standardized data, first for non-overlapping percentiles of 5 percent each, then for cumulative percentiles from 5 percent to 50 percent coverage. Figure 5 matches these empirical estimates to the theoretical truncated correlations for the standardized S&P500 and FTSE100 returns. Note that the theoretical truncated correlation function is based on a bivariate Student- t distribution with $\hat{r} = 4$ joint degrees of freedom (see Table 3). Figure 5 also gives the simulated 95% confidence limits.

INSERT FIGURE 5

The non-overlapping empirical results do seem to follow the general U-shape postulated in Section 2 and indicated by the theoretical function in Figure 5, panel A. It would appear that the variability in the empirical truncated correlations is well within the 95 percent confidence interval. The cumulative empirical truncated correlation estimates in Figure 5, panel B, also fit within the 95% confidence interval. Interestingly, the empirical inverted U-shape would better fit a theoretical Student- t distribution with a larger joint degrees of freedom parameter than the empirically estimated $\hat{r} = 4$ (cf., Figure 3). In any case, the empirical evidence does not support size-dependent correlation or a meltdown in diversification.

Of course, the evidence has to be rather strong before we are able to reject the size-independent correlation null hypothesis. The standard error function in Figure 4 suggests that we need a large sample size to obtain a sufficiently narrow confidence interval.

To compare a low (cross-Atlantic) unconditional correlation pair with a high (inter-European) unconditional correlation pair, we also present and discuss the FTSE100 and

CAC40 combination with $\hat{\rho} = 0.65$ ¹⁰. The results in Figure 6 consist of four panels. Panels A and B illustrate the truncated correlation estimates for non-overlapping percentiles. Panels C and D illustrate the truncated correlation estimates for cumulative percentiles. Panels A and C are based on a bivariate Student- t distribution with $\hat{r} = 5$ joint degrees of freedom. Panels B and D are based on a normal distribution.

INSERT FIGURE 6

Under the assumption of normality, the empirical results (Panels B and D) would lead us to conclude that truncated correlation between standardized FTSE100 and CAC40 returns increases for large movements in the FTSE100 returns. In fact, panels B and D suggest a significant diversification meltdown for extremely negative returns on the FTSE100, since the lowest left tail percentile estimate in panel B is outside the 95 percent confidence limits. The cumulative impact of that single exceedance is evident in panel D. However, this size-dependent increase in correlation may simply be due to the fact that the returns are better parameterised by a bivariate Student- t distribution. Table 3 suggests that the FTSE100 and CAC40 return distribution is best characterized by a Student- t with $\hat{r} = 5$ degrees of freedom. Panel A, Figure 6, indicates that the relatively high empirical truncated correlation estimates almost perfectly fit the theoretical truncated correlation function based on this particular Student- t assumption. The standard errors in Figure 6 (panels A and C) indicate that for the left-tail percentiles, we cannot reject the null hypothesis of size-independent correlation for large negative FTSE100 returns. Surprisingly, the results do suggest that the upper right tail percentile estimate is outside (below) its 95 percent confidence limits. This negative size-dependency suggests a decrease (increase) in correlation (diversification). Of course, this happens exactly when it is least wanted by investors who are long in British and French

¹⁰ The graphs for other combinations of return series are available from the authors on request.

stocks. This apparent asymmetry in truncated correlations (not unlike the findings in Ang and Chen, 2002) could indicate a mixture of Student- t distributions with e.g., respectively $\hat{\nu} = 4$ (for negative returns on the FTSE100) and $\hat{\nu} = 6$ (for positive returns on the FTSE100) degrees of freedom. Please note that we imposed a single degrees of freedom parameter on our bivariate Student- t distribution. Nevertheless, even if right tail $\hat{\nu}$ were to increase to its limiting value, we would still reject the null of size-independent correlation for large positive FTSE100 returns. In that case, the normal distribution applies and Panel D, Figure 6, indicates that the right-tail percentiles are still outside (below) the confidence bands. We found similar evidence of size-dependent correlation for large positive returns in the DAX-S&P500, and DAX-FTSE100 combinations. In both cases, our findings suggest an increase in correlation. For the S&P500-USGB, we found evidence of size-dependent correlation for large negative returns in the S&P500. Here, also, the size-dependency indicated an increase in correlation. None of our results therefore support the diversification meltdown hypothesis.

4. Portfolio allocation implications of size-dependent correlation

The previous sections illustrate the risks involved in drawing conclusions about a breakdown in diversification based on an incorrect distributional assumption. The empirical Section 3 provides a more careful approach to determine whether there is evidence supporting size-conditional correlation. We did find limited support for size-conditional correlation, but it was not in the expected direction. Nevertheless, it seems worthwhile to investigate how size-conditionality could possibly affect finance practitioners. In Section 2, we mentioned the relevance for beta estimation. We generalize the discussion here, to include the implications of size-dependent correlation on portfolio allocation¹¹.

¹¹ Loretan and English (2000a) discuss the impact of increasing conditional correlation on portfolio Value-at-Risk (VaR). Individual asset VaRs are often computed based on short observation periods to correctly reflect VaR sensitivity to conditionality in variance. Portfolio VaRs are similarly sensitive to conditionality in correlation.

The benefits of international portfolio diversification depend crucially on asset returns being less than perfectly correlated. These benefits could be overrated if correlation is found to increase in the lower tails of the bivariate return distributions. In fact, when most needed, the protection offered by diversification would rapidly vanish. This finding necessitates a revision of the mean-variance portfolio allocation model. For short-horizon investors, this need not necessarily be a problem. Their dynamic portfolio allocation would be based on the conditional correlation matrix, which would be their best forecast correlation matrix if there is persistence in the size of returns (e.g., a persistent bear market). That is, they would immediately capture the excessive correlation and adjust their portfolios accordingly. A straightforward implementation of short-horizon dynamic asset allocation is given by Turtle, Buse and Korkie (1994). For long-horizon investors, the issue is more complicated. Their best forecast correlation matrix would be based on the unconditional correlation matrix (since bear markets do not persist indefinitely). If there is still evidence of size-dependent correlation, after filtering for intertemporal dependency, then this would suggest a flaw in standard mean-variance optimization. Instead of maximising expected returns given the unconditional correlation matrix, long-term investors would then have to maximise expected returns given a size-conditional correlation matrix. That would probably require higher-order moments (in addition to mean and variance) to also be included in the investor's utility function.

Our results in Section 3 suggest that this could in fact be necessary. After inferring the unconditional correlation from the biased conditional correlation estimates, we find significant evidence of size-dependent correlation. This suggests that adjustments to the portfolio allocation process will be necessary. Even without size-dependent correlation, certain adjustments to the allocation process may still be necessary due to the apparent non-normality of the data.

5. Conclusions

Turbulent financial market conditions easily lead to an impression of contagion and spillover effects wreaking havoc on the benefits of international diversification. The obvious conclusion is that correlation between international financial asset returns increases with the size of (usually negative) returns. If this intuition is corroborated by empirical evidence, it would have serious implications for international portfolio allocation, in particular for short-horizon investors given the persistence in the size of returns. A variety of conditional correlation estimators have recently been proposed to measure this diversification meltdown effect. Unfortunately, most of these estimators suffer to some extent from estimation bias making comparisons difficult or even impossible. In this paper, we evaluate the performance of one popular conditional correlation estimator, the truncated correlation estimator that conditions on non-overlapping and/or cumulative percentiles of the bivariate return distribution. It is relatively straightforward to capture the estimation bias for this truncated correlation estimator under the assumption of conditional normality. Since the joint conditional normality assumption is unlikely to be valid for financial asset returns, we analytically derive and measure the bias in this estimator for fat-tailed bivariate Student- t distributions. The estimation bias is now considerably larger than for the bivariate normal distribution. This suggests that earlier studies may have overestimated the size-dependency in correlation, simply due to the assumption of bivariate normality.

When applied to a data set of standardized international stock market index returns, we find that, under the assumption of normally distributed returns, there is evidence of (positive) size-dependent correlation. This would indicate that a size-conditional variance-covariance matrix ought to be used for dynamic mean-variance portfolio allocation. However, when assuming the more likely Student- t distribution, we find that positive size-dependency

disappears. Instead, we find significant evidence of negative size-dependency, indicating a strengthening of diversification for large positive returns. After measuring the Student- t bias in the truncated correlation estimate, from short, select, samples of historical returns, we can then infer the relevant correlation estimate to be used in the portfolio allocation. Note, however, that the allocation exercise should also account for the non-normality of the joint return distributions.

References

- Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. *Journal of Financial Economics* 63, 443-494.
- Ball, C.A., Torous, W.N., 2000. Stochastic correlation across international stock Markets. Working Paper # 17-00, UCLA.
- 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., Engle, R.F., and 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 Nr. 597R, Board of Governors of the Federal Reserve Board: Washington, D.C.
- Butler, K.C., Joaquin, 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.A., Koedijk, C.G., Kofman, P., 2002. Increased correlation in bear markets. *Financial Analysts Journal* 58 (1), 87-94.
- 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 & Economic Statistics* 20, 339-350.
- Engle, R.F., Lin, W.-L., Ito, T., 1994. Do bulls and bears move across borders? International transmission of stock returns and volatility. *Review of Financial Studies* 7, 507-538.
- Engle, R.F., Kroner, K., 1995. Multivariate simultaneous GARCH. *Econometric Theory* 11, 122-150.
- Forbes, K.J., Rigobon, R., 2002. No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance* 57, 2223-2261.
- Hill, D.M., 1975. A simple approach to inference about the tail of a distribution. *Annals of Statistics* 3, 1163-1174.
- Huisman, R., Koedijk, K.G., Pownall, R.A., 1998. VaR-x: Fat tails in financial risk management. *Journal of Risk* 1, 43-61.

- Johnson, N.L., Kotz, S., 1972. Distributions in statistics: Continuous multivariate distributions. John Wiley & Sons, Inc.
- 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.
- Kritzman, M., Lowry, K., Van Royen, A-S., 2001. Risk, regimes, and overconfidence. *Journal of Derivatives* 8, 32- 42.
- 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 Nr. 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*, June, 29-36.
- McDonald, I., 2000. The myth of going overseas for diversification. *TheStreet.com*.
- Patton, A., 2001. Modelling time-varying exchange rate dependence using the conditional copula. University of California, San Diego, Discussion Paper 01-09.
- Regier, M.H., Hamdan, M.A., 1971. Correlation in a bivariate normal distribution with truncation in both variables. *Australian Journal of Statistics* 13, 77-82.
- Sullivan, G., 1995. Correlation counts. *RISK Magazine* 8, 36-38.
- Turtle, H., Buse, A., Korkie, B., 1994. Tests of conditional asset pricing with time-varying moments and risk prices. *Journal of Financial and Quantitative Analysis* 29, 15-29.

Appendix – Mean and Variance of the Truncated Student- t

Let T have a Student- t distribution with r degrees of freedom and denote

$$P_{[L,U]}^r = \Pr(L \leq T \leq U) = \int_L^U f_{t_r}(t) dt \quad (\text{A1})$$

where

$$f_{t_r}(t) = C_r \left(1 + \frac{t^2}{r}\right)^{-(r+1)/2} \quad \text{and} \quad C_r = \frac{\Gamma((r+1)/2)}{(r\pi)^{1/2} \Gamma(r/2)}. \quad (\text{A2})$$

Now let S have a truncated Student- t distribution with r degrees of freedom, so that

$$\Pr(L \leq S \leq U) = \frac{1}{P_{[L,U]}^r} \int_L^U f_{t_r}(t) dt = 1. \quad (\text{A3})$$

Then

$$E[S] = \frac{C_r}{P_{[L,U]}^r} \int_L^U s \left(1 + \frac{s^2}{r}\right)^{-(r+1)/2} ds = \frac{rC_r}{(r-1)P_{[L,U]}^r} \left[\left(1 + \frac{L^2}{r}\right)^{-(r-1)/2} - \left(1 + \frac{U^2}{r}\right)^{-(r-1)/2} \right] \quad (\text{A4})$$

The variance of S can be determined from the usual relationship

$$\text{Var}[S] = E[S^2] - E[S]^2, \quad (\text{A5})$$

where for the truncated Student- t_r distribution,

$$E[S^2] = \frac{rC_r}{(r-1)P_{[L,U]}^r} \left[L \left(1 + \frac{L^2}{r}\right)^{-(r-1)/2} - U \left(1 + \frac{U^2}{r}\right)^{-(r-1)/2} + \frac{(\sqrt{r})P_{[L^*,U^*]}^{r-2}}{(\sqrt{r}-2)C_{r-2}} \right] \quad (\text{A6})$$

$$\text{with} \quad L^* = \frac{L\sqrt{r-2}}{\sqrt{r}} \quad \text{and} \quad U^* = \frac{U\sqrt{r-2}}{\sqrt{r}}.$$

To derive this result, note that

$$E[S^2] = \frac{C_r}{P_{[L,U]}^r} \int_L^U s^2 \left(1 + \frac{s^2}{r}\right)^{-(r+1)/2} ds. \quad (\text{A7})$$

The integral can be solved using *integration by parts*, i.e.

Let

$$u = \frac{rs}{2} \quad dv = \frac{2s}{r} \left(1 + \frac{s^2}{r}\right)^{-(r+1)/2} ds$$

so

$$du = \frac{r}{2} ds \quad v = \frac{\left(1 + \frac{s^2}{r}\right)^{-(r-1)/2}}{-(r-1)/2}$$

and

$$\begin{aligned} E[s^2] &= \frac{C_r}{P_{[L,U]}^r} \left[\frac{rs \left(1 + \frac{s^2}{r}\right)^{-(r-1)/2}}{-(r-1)} \Big|_L^U + \frac{r}{(r-1)} \int_L^U \left(1 + \frac{s^2}{r}\right)^{-(r-1)/2} ds \right] \\ &= \frac{rC_r}{(r-1)P_{[L,U]}^r} \left[L \left(1 + \frac{L^2}{r}\right)^{-(r-1)/2} - U \left(1 + \frac{U^2}{r}\right)^{-(r-1)/2} + \int_L^U \left(1 + \frac{s^2}{r}\right)^{-(r-1)/2} ds \right] \end{aligned} \quad (\text{A8})$$

The remaining integral on the right hand side of (A8) is related to probabilities associated with an untruncated Student- t distribution with $(r-2)$ degrees of freedom. In particular,

$$\int_L^U \left(1 + \frac{s^2}{r}\right)^{-(r-1)/2} ds = \frac{\sqrt{r}}{\sqrt{r-2}} \int_{L^*}^{U^*} \left(1 + \frac{w^2}{a}\right)^{-(a+1)/2} dw = \frac{(\sqrt{r})P_{[L^*,U^*]}^{r-2}}{(\sqrt{r-2})C_{r-2}}, \quad (\text{A9})$$

and hence the result in (A6).

Table 1. Summary statistics for index returns

	S&P 500	FTSE 100	CAC 40	DAX 100	USGB 10yr
<u>Descriptive statistics for raw returns</u>					
Annualized Mean	16.52%	13.43%	12.84%	10.82%	6.81%
Annualized Standard Deviation	13.93%	12.62%	18.91%	18.07%	6.37%
Maximum Daily Return	4.99%	5.70%	6.81%	6.44%	1.67%
Minimum Daily Return	-7.11%	-3.54%	-7.57%	-10.05%	-2.83%
Skewness	-0.35*	0.08	-0.16*	-0.64*	-0.39*
Kurtosis	5.55*	3.08	2.58*	6.72*	3.00
Jarque-Bera test statistic**	756.32*	3.40	29.51*	1665.48*	65.30*
<u>Tail statistics for standardized returns</u>					
Hill upper tail index α***	5.00	4.23	5.43	4.15	5.04
Hill lower tail index α***	4.59	3.77	4.24	3.76	4.56
t-degrees of freedom r	3.51	4.74	4.99	3.55	4.82

The table gives the summary statistics for daily (standardized) return indices: S&P 500 Composite Index, FTSE 100 All Share Index, CAC 40 Index, DAX 100 Performance Index and the 10-Year US Benchmark Government Bond Index over the period January 1990 - December 1999 ($N=2580$ daily observations).

* indicates significantly different from normal distribution values at 95% confidence level.

** Jarque Bera normality test: $JB = N \left[\frac{skewness^2}{6} + \frac{(kurtosis - 3)^2}{24} \right] \sim \chi_2^2$.

*** Hill tail index estimator: $\hat{\alpha} = \left[\frac{1}{m-1} \sum_{i=1}^{m-1} \ln(X_{(i)}) - \ln(X_{(m)}) \right]^{-1}$ where the $X_{(i)}$ are the descending order statistics of the returns X and m is selected (pragmatically) at 2% of the total sample size N (here: 258 observations).

• **Table 2. Unconditional correlation matrix for index returns**

Panel A – Raw Returns

	S&P 500	FTSE 100	CAC 40	DAX 100	USGB 10yr
S&P 500	1				
FTSE 100	0.350	1			
CAC 40	0.367	0.661	1		
DAX 100	0.294	0.573	0.623	1	
USGB 10yr	0.273	0.103	0.124	0.037	1

Panel B – Standardized Returns

	S&P 500	FTSE 100	CAC 40	DAX 100	USGB 10yr
S&P 500	1				
FTSE 100	0.345	1			
CAC 40	0.331	0.651	1		
DAX 100	0.249	0.524	0.597	1	
USGB 10yr	0.340	0.141	0.147	0.044	1

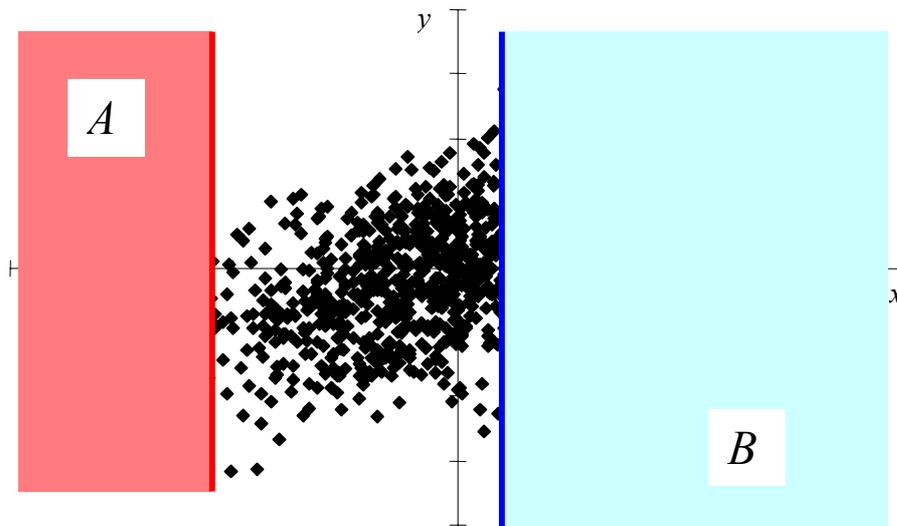
The table gives the unconditional correlation matrix for the sample period January 1990 - December 1999 for the following daily index return series: S&P 500 Composite Index, FTSE 100 All Share Index, CAC 40 Index, DAX 100 Performance Index and the 10-Year US Benchmark Government Bond Index.

Table 3. Bivariate Student-*t* joint degrees of freedom parameter estimates

	S&P 500	FTSE 100	CAC 40	DAX 100	USGB 10yr
S&P 500	3.507				
FTSE 100	3.970	4.741			
CAC 40	4.420	5.060	4.994		
DAX 100	3.640	3.780	4.030	3.546	
USGB 10yr	4.850	6.200	6.320	4.490	4.824

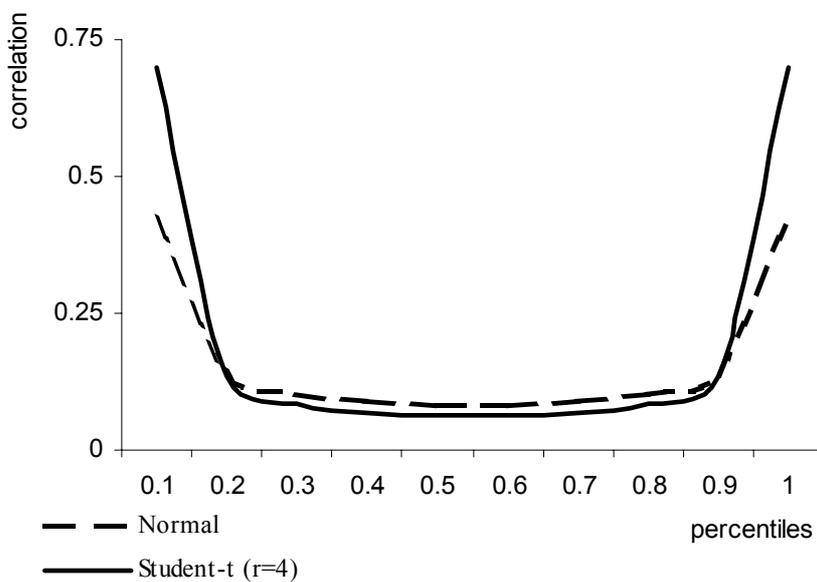
The table gives the joint *t*-degrees of freedom parameter, $\hat{\nu}$, estimates for the bivariate Student-*t* distribution for the standardized index return series using Maximum Likelihood Estimation.

Figure 1. Marginal conditioning of the joint returns distribution



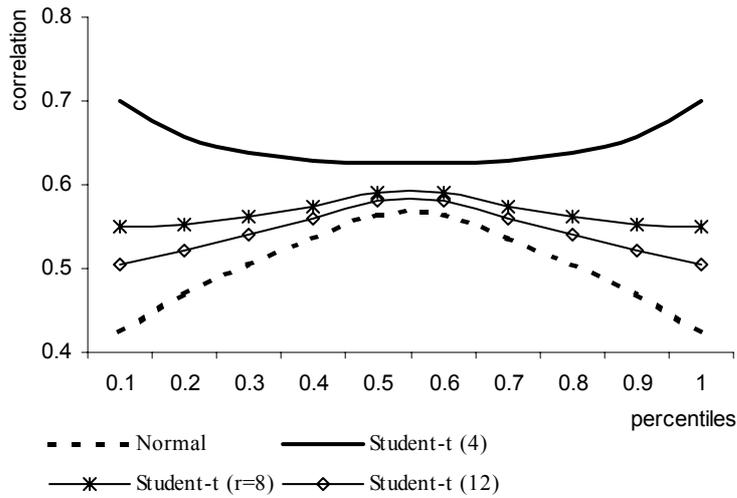
Two possible partitionings of a bivariate normal distribution with unconditional correlation $\rho = 0.5$. Conditioning occurs on one marginal component, x .

Figure 2. Theoretical truncated correlation functions (non-overlapping percentiles)



The figure gives theoretical truncated correlations for ten deciles of the joint distribution of returns. We assume unconditional correlation $\rho=0.75$ and compare the theoretical truncated correlations for a bivariate normal with those for a bivariate Student- t distribution ($r=4$ degrees of freedom).

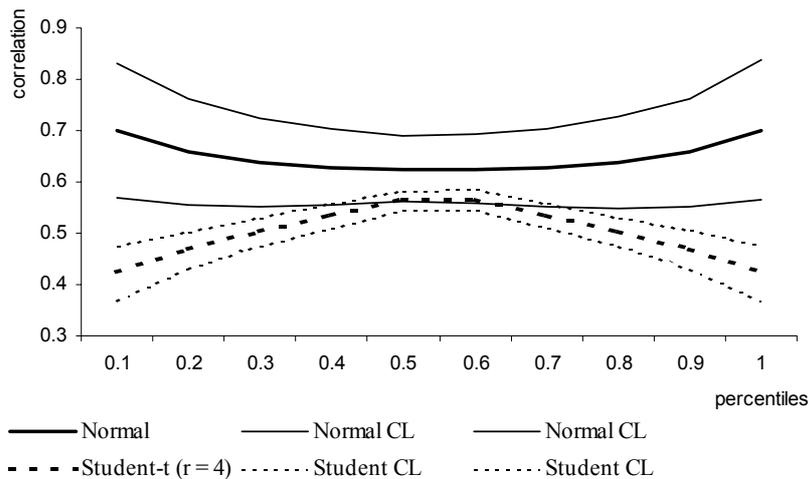
Figure 3. Theoretical truncated correlation functions (cumulative percentiles)



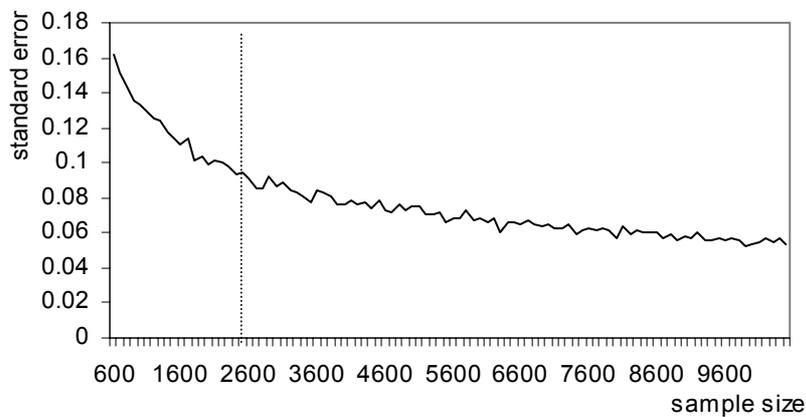
The figure gives theoretical truncated correlations for cumulative percentiles of the joint distribution of returns. We assume unconditional correlation $\rho=0.75$, and compare the theoretical truncated correlations for a bivariate normal with those for three bivariate Student- t distributions ($r=4, 8, 12$ degrees of freedom).

Figure 4. Simulated standard errors: precision and sample size

Panel A

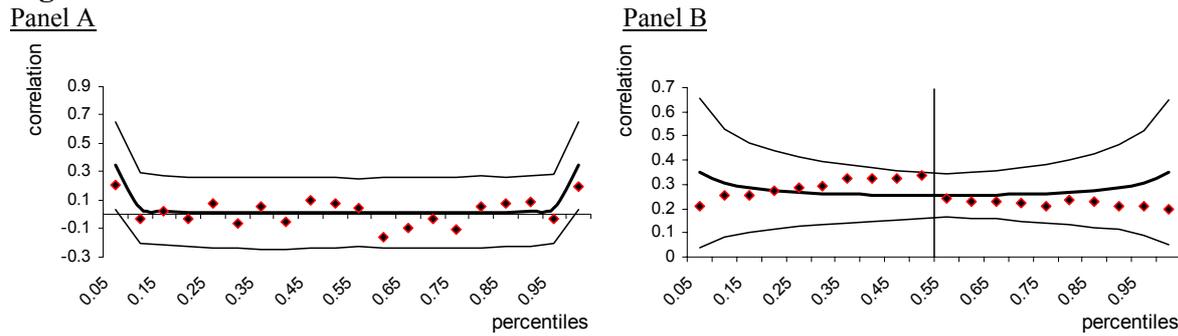


Panel B



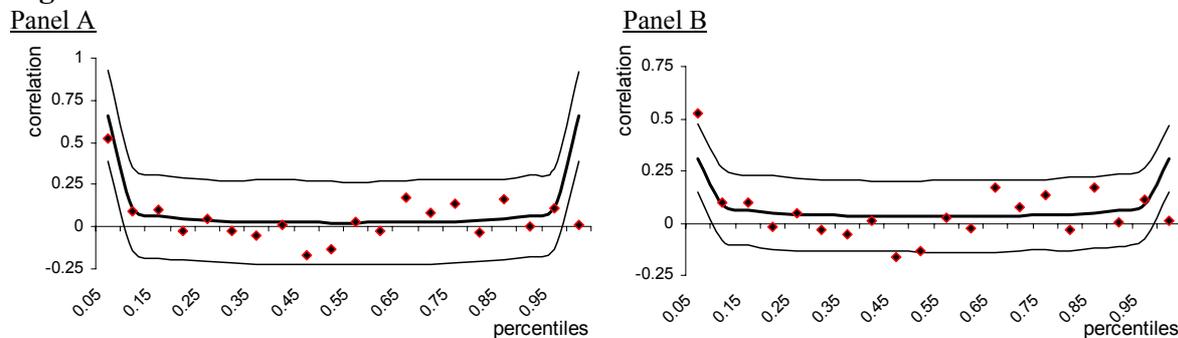
Standard errors in panel B are computed by simulation for the 5 percent left tail truncated correlation.

Figure 5. S&P500 versus FTSE100 truncated correlation

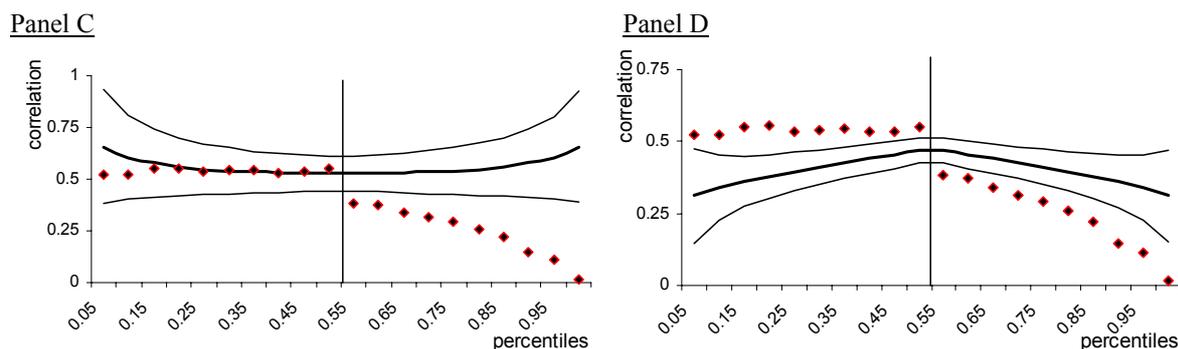


Truncated correlations are computed for fixed percentiles (panel A) and cumulative percentiles (panel B). The solid black line illustrates the theoretical truncated correlation function for a bivariate Student- t distribution with $\nu=4$ degrees of freedom, and unconditional correlation $\rho=0.35$. The thin black lines indicate the 95% simulated confidence limits. The diamonds indicate the empirical truncated correlations. The vertical black line indicates the cumulative sample split between right and left tail returns

Figure 6. FTSE100 versus CAC40 truncated correlation



Panel A assumes a Student- t ($\nu=5$) distribution; Panel B assumes a normal distribution. Non-overlapping truncated correlations are computed for fixed 2.5 percent percentiles. The solid black line illustrates the theoretical truncated correlation function for a bivariate Student- t distribution with $\nu=5$ degrees of freedom, and unconditional correlation $\rho=0.65$. The thin black lines indicate 95% simulated confidence limits. The diamonds indicate the empirical truncated correlations.



Panel C assumes a Student- t ($df=5$) distribution; Panel D assumes a normal distribution. Cumulative truncated correlations are computed for increasing percentiles starting in the outer left tail up to the median [0 to 0.50], respectively the outer right tail up to the median [1 to 0.50]. The solid black line illustrates the theoretical truncated correlation function for a bivariate Student- t distribution with $\nu=5$ degrees of freedom, and unconditional correlation $\rho=0.65$. The thin black lines indicate 95% simulated confidence limits. The diamonds indicate the empirical truncated correlations. The vertical black lines indicate the cumulative sample split between right and left tail returns.