

International Journal of Theoretical and Applied Finance
© World Scientific Publishing Company

THE JOINT DISTRIBUTION OF STOCK RETURNS IS NOT ELLIPTICAL

RÉMY CHICHEPORTICHE

*Capital Fund Management, 6–8 boulevard Haussmann
Paris, 75009, France
Remy.Chicheportiche@cfm.fr*

JEAN-PHILIPPE BOUCHAUD

*Capital Fund Management, 6–8 boulevard Haussmann
Paris, 75009, France
Jean-Philippe.Bouchaud@cfm.fr*

Received (Day Month Year)

Revised (Day Month Year)

Using a large set of daily US and Japanese stock returns, we test in detail the relevance of Student models, and of more general elliptical models, for describing the joint distribution of returns. We find that while Student copulas provide a good approximation for strongly correlated pairs of stocks, systematic discrepancies appear as the linear correlation between stocks decreases, that rule out *all* elliptical models. Intuitively, the failure of elliptical models can be traced to the inadequacy of the assumption of a single volatility mode for all stocks. We suggest several ideas of methodological interest to efficiently visualise and compare different copulas. We identify the rescaled difference with the Gaussian copula and the central value of the copula as strongly discriminating observables. We insist on the need to shun away from formal choices of copulas with no financial interpretation.

Keywords: Copulas; Stock returns; Multivariate distribution; Linear correlation; Non-linear dependences; Student distribution; Elliptical distributions

1. Introduction

The most important input of portfolio risk analysis and portfolio optimisation is the correlation matrix of the different assets. In order to diversify away the risk, one must avoid allocating on bundles of correlated assets, an information in principle contained in the correlation matrix. The seemingly simple mean-variance Markowitz program is however well known to be full of thorns. In particular, the empirical determination of large correlation matrices turns out to be difficult, and some astute “cleaning” schemes must be devised before using it for constructing optimal allocations [26]. This topic has been the focus of intense research in the recent years, some inspired by Random Matrix Theory (for a review see [13,38]) or clustering ideas [32,42,43].

There are however many situations of practical interest where the (linear) correlation matrix is inadequate or insufficient [1,30,33]. For example, one could be more interested in minimising the probability of large negative returns rather than the variance of the portfolio. Another example is the Gamma-risk of option portfolios, where the correlation of the squared-returns of the underlyings is needed. Credit derivatives are bets on the probability of simultaneous default of companies or of individuals; again, an estimate of the correlation of tail events is required (but probably very hard to ascertain empirically) [19], and for a recent interesting review [2].

Apart from the case of multivariate Gaussian variables, the description of non-linear dependence is not reducible to the linear correlation matrix. The general problem of parameterising the full joint probability distribution of N random variables can be “factorised” into the specification of all the marginals on the one hand, and of the dependence structure (called the ‘copula’) of N standardised variables with uniform distribution in $[0, 1]$, on the other hand. The nearly trivial statement that all multivariate distributions can be represented in that way is called Sklar’s Theorem [14,41]. Following a typical pattern of mathematical finance, the introduction of copulas ten years ago has been followed by a calibration spree, with academics and financial engineers frantically looking for copulas to best represent their pet multivariate problem. But instead of trying to understand the economic or financial mechanisms that lead to some particular dependence between assets and construct copulas that encode these mechanisms, the methodology has been — as is sadly often the case — to brute force calibrate on data copulas straight out from statistics handbooks [12,15,16,17,36]. The “best” copula is then decided from some quality-of-fit criterion, irrespective of whether the copula makes any intuitive sense (some examples are given below). This is reminiscent of the ‘local volatility models’ for option markets [11]: although the model makes no intuitive sense and cannot describe the actual dynamics of the underlying asset, it is versatile enough to allow the calibration of almost any option smile (see [22]). This explains why this model is heavily used in the financial industry. Unfortunately, a blind calibration of some unwarranted model (even when the fit is perfect) is a recipe for disaster. If the underlying reality is not captured by the model, it will most likely derail in rough times — a particularly bad feature for risk management! Another way to express our point is to use a Bayesian language: there are families of models for which the ‘prior’ likelihood is clearly extremely small — we discuss below the case of Archimedean copulas. Statistical tests are not enough — intuition and plausibility are mandatory.

The aim of this paper is to study in depth the family of elliptical copulas, in particular Student copulas, that have been much used in a financial context and indeed have a simple intuitive interpretation [14,18,23,27,30,40]. We investigate in detail whether or not such copulas can faithfully represent the joint distribution of the returns of US stocks. (We have also studied other markets as well). We unveil

clear, systematic discrepancies between our empirical data and the corresponding predictions of elliptical models. These deviations are qualitatively the same for different tranches of market capitalisations, different time periods and different markets. Based on the financial interpretation of elliptical copulas, we argue that such discrepancies are actually expected, and propose the ingredients of a generalisation of elliptical copulas that should capture adequately non-linear dependences in stock markets. The full study of this generalised model, together with a calibration procedure and stability tests, will be provided in a subsequent publication.

The outline of the paper is as follows. We first review different measures of dependence (non-linear correlations, tail dependence), including copulas. We discuss simple ways to visualise the information contained in a given copula. We then introduce the family of elliptical copulas, with special emphasis on Student copulas and log-normal copulas, insisting on their transparent financial interpretation, in contrast with the popular but rather implausible Archimedean copulas. We then carefully compare our comprehensive empirical data on stock markets with the predictions of elliptical models — some of which being parameter free, and conclude that elliptical copulas fail to describe the full dependence structure of equity markets.

2. Bivariate measures of dependence and copulas

In this section, we recall several bivariate measures of dependence, and study their properties when the distribution of the random pair under scrutiny is elliptical. The incentive to focus on bivariate measures comes from the theoretical property that all the marginals, including bivariate, of a multivariate elliptical distribution are themselves elliptical. In turn, a motivated statement that the pairwise distributions are not elliptical is enough to claim the non-ellipticity of the joint multivariate distribution. This is the basic line of argumentation of the present paper.

2.1. Correlation coefficients

Beyond the standard correlation coefficient, one can characterise the dependence structure through the correlation between powers of the random variables of interest:

$$\rho_{ij}^{(d)} = \frac{\mathbb{E}[S_i S_j |X_i X_j|^d] - \mathbb{E}[S_i |X_i|^d] \mathbb{E}[S_j |X_j|^d]}{\sqrt{\mathbb{V}[S_i |X_i|^d] \mathbb{V}[S_j |X_j|^d]}} \quad (2.1)$$

$$\zeta_{ij}^{(d)} = \frac{\mathbb{E}[|X_i X_j|^d] - \mathbb{E}[|X_i|^d] \mathbb{E}[|X_j|^d]}{\sqrt{\mathbb{V}[|X_i|^d] \mathbb{V}[|X_j|^d]}} \quad (2.2)$$

where $S_i = \text{sign}(X_i)$ and provided the variances \mathbb{V} in the denominators are well defined. The case $\rho^{(d)}$ for $d = 1$ corresponds to the usual linear correlation coefficient, for which we will use the standard notation ρ , whereas $\zeta^{(d)}$ for $d = 2$ is the correlation of the squared returns, that would appear in the Gamma-risk of a Δ -hedged option

portfolio. However, high values of d are expected to be very noisy in the presence of power-law tails (as is the case for financial returns) and one should seek for lower order moments, such as $\zeta^{(d)}$ for $d = 1$ which also captures the correlation between the *amplitudes* (or the volatility) of price moves, or even $\rho^{(s)} \equiv \rho^{(d)}$ for $d = 0$ that measures the correlation of the signs.

In the case of bivariate Gaussian variables with zero mean, all these correlations can be expressed in terms of ρ . For example, the quadratic correlation $\zeta^{(2)} = \zeta^{(d)}$ for $d = 2$ is given by:

$$\zeta^{(2)} = \rho^2; \quad (2.3)$$

whereas the correlation of absolute values $\zeta^{(1)} = \zeta^{(d)}$ for $d = 1$ is given by:

$$\zeta^{(1)} = \frac{D(\rho) - 1}{\frac{\pi}{2} - 1}, \quad (2.4)$$

with $D(r) = \sqrt{1-r^2} + r \arcsin r$. The correlation of signs is given by $\rho^{(s)} = \frac{2}{\pi} \arcsin \rho$. For some other classes of distributions, the higher-order coefficients $\rho^{(d)}$ and $\zeta^{(d)}$ are explicit functions of the coefficient of linear correlation. This is for example the case of Student variables (see Fig. 1) and more generally for all elliptical distributions (see below).

2.2. Tail dependence

Another characterisation of dependence, of great importance for risk management purposes, is the so-called *coefficient of tail dependence* which measures the joint probability of extreme events. More precisely, the upper tail dependence is defined as [14,30]:

$$\tau_{UU}(p) = \mathbb{P} [X_i > \mathcal{P}_{<,i}^{-1}(p) \mid X_j > \mathcal{P}_{<,j}^{-1}(p)], \quad (2.5)$$

where $\mathcal{P}_{<,k}$ is the cumulative distribution function (cdf) of X_k , and p a certain probability level. In spite of its seemingly asymmetric definition, it is easy to show that τ_{UU} is in fact symmetric in $X_i \leftrightarrow X_j$ when all univariate marginals are identical. When $p \rightarrow 1$, $\tau_{UU}^* = \tau_{UU}(1)$ measures the probability that X_i takes a very large positive value knowing that X_j is also very large, and defines the asymptotic tail dependence. Random variables can be strongly dependent from the point of view of linear correlations, while being nearly independent in the extremes. For example, bivariate Gaussian variables are such that $\tau_{UU}^* = 0$ for any value of $\rho < 1$. The lower tail dependence τ_{LL} is defined similarly:

$$\tau_{LL}(p) = \mathbb{P} [X_i < \mathcal{P}_{<,i}^{-1}(1-p) \mid X_j < \mathcal{P}_{<,j}^{-1}(1-p)], \quad (2.6)$$

and is equal to $\tau_{UU}(p)$ for symmetric bivariate distributions. One can also define mixed tail dependence, for example:

$$\tau_{LU}(p) = \mathbb{P} [X_i < \mathcal{P}_{<,i}^{-1}(1-p) \mid X_j > \mathcal{P}_{<,j}^{-1}(p)], \quad (2.7)$$

with obvious interpretations.

2.3. Copulas

There are many other possible coefficients of dependence, such as Spearman's rho or Kendall's tau, that measure, respectively, the correlation of ranks or the "concordance" probability (see e.g. [39] for an introduction). Both these measures are invariant under any increasing transformations. More generally, the copula (generalisable to dimensions larger than 2), encodes all the dependence between random variables that is invariant under increasing transformations. The copula of a random pair (X_i, X_j) is the joint cdf of $U_i = \mathcal{P}_{<,i}(X_i)$ and $U_j = \mathcal{P}_{<,j}(X_j)$:

$$C(u_i, u_j) = \mathbb{P}[\mathcal{P}_{<,i}(X_i) \leq u_i \text{ and } \mathcal{P}_{<,j}(X_j) \leq u_j] \quad (2.8)$$

Since the marginals of U_i and U_j are uniform by construction, the copula only captures their degree of "entanglement". For independent random variables, $C(u_i, u_j) \equiv u_i u_j$.

Whereas the $\rho^{(d)}$'s and $\zeta^{(d)}$'s depend on the marginal distributions, the tail dependences, Spearman's rho and Kendall's tau can be fully expressed in terms of the copula only. For example, Kendall's tau is given by $\tau^{(K)} = 4 \int_{[0,1]^2} C(u_i, u_j) dC(u_i, u_j) - 1$, while the tail dependence coefficients can be expressed as:

$$\frac{1 - 2p + C(p, p)}{1 - p} = \tau_{UU}(p) \quad \frac{C(1-p, 1-p)}{1 - p} = \tau_{LL}(p) \quad (2.9a)$$

$$\frac{1 - p - C(p, 1-p)}{1 - p} = \tau_{UL}(p) \quad \frac{1 - p - C(1-p, p)}{1 - p} = \tau_{LU}(p). \quad (2.9b)$$

Copulas are not easy to visualise, first because they need to be represented as 3-D plot of a surface in two dimensions and second because there are bounds (called Fréchet bounds [30]) within which $C(u, v)$ is constrained to live, and that compress the difference between different copulas (the situation is even worse in higher dimensions). Estimating copula densities, on the other hand, is even more difficult than estimating univariate densities, especially in the tails. We therefore propose to focus on the diagonal of the copula, $C(p, p)$ and the anti-diagonal, $C(p, 1-p)$, that capture part of the information contained in the full copula, and can be represented as 1-dimensional plots. Furthermore, in order to emphasise the difference with the case of independent variables, it is expedient to consider the following quantities:

$$\frac{C(p, p) - p^2}{p(1-p)} = \tau_{UU}(p) + \tau_{LL}(1-p) - 1 \quad (2.10a)$$

$$\frac{C(p, 1-p) - p(1-p)}{p(1-p)} = 1 - \tau_{UL}(p) - \tau_{LU}(1-p), \quad (2.10b)$$

where the normalisation is chosen such that the tail correlations appear naturally. Note in particular that the diagonal quantity tends to the asymptotic tail dependence coefficients τ_{UU}^* (τ_{LL}^*) when $p \rightarrow 1$ ($p \rightarrow 0$). Similarly, the anti-diagonal tends to $-\tau_{UL}^*$ as $p \rightarrow 1$ and to $-\tau_{LU}^*$ as $p \rightarrow 0$.

Another important reference point is the Gaussian copula $C_G(u, v)$, and we will consider below the normalised differences along the diagonal and the anti-diagonal:

$$\Delta_d(p) = \frac{C(p, p) - C_G(p, p)}{p(1-p)}; \quad \Delta_a(p) = \frac{C(p, 1-p) - C_G(p, 1-p)}{p(1-p)}, \quad (2.11)$$

which again tend to the asymptotic tail dependence coefficients in the limits $p \rightarrow 1$ ($p \rightarrow 0$), owing to the fact that these coefficients are all zero in the Gaussian case.

The centerpoint of the copula, $C(\frac{1}{2}, \frac{1}{2})$, is particularly interesting: it is the probability that both variables are simultaneously below their respective median. For bivariate Gaussian variables, one can show that:

$$C_G(\frac{1}{2}, \frac{1}{2}) = \frac{1}{4} + \frac{1}{2\pi} \arcsin \rho \equiv \frac{1}{4} (1 + \rho^{(s)}), \quad (2.12)$$

where $\rho^{(s)}$ is the sign correlation coefficient defined above. The trivial but remarkable fact is that the above expression holds for any elliptical models, that we define and discuss in the next section. This will provide a powerful test to discriminate between empirical data and a large class of copulas that have been proposed in the literature.

3. Review of elliptical models

3.1. General elliptical models

Elliptical random variables X_i can be simply generated by multiplying standardised Gaussian random variables ϵ_i with a common random (strictly positive) factor σ , drawn independently from the ϵ 's from an arbitrary distribution $P(\sigma)$: [14]

$$X_i = \mu_i + \sigma \cdot \epsilon_i \quad i = 1 \dots N \quad (3.1)$$

where μ_i is a location parameter (set to zero in the following), and $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$. Such models are called ‘‘elliptical’’ because the corresponding multivariate probability distribution function (pdf) depends only on the rotationally invariant quadratic form $(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})$; therefore quantile levels define ellipses in N dimensions (see [4] for the construction and properties of elliptically contoured multivariate distributions).

In a financial context, elliptical models are basically stochastic volatility models with arbitrary dynamics such that the ergodic distribution is $P(\sigma)$ (since we model only single-time distributions, the time ordering is irrelevant here). The important assumption, however, is that the random amplitude factor σ is *equal for all stocks* (or all assets in a more general context). In other words, one assumes that the mechanisms leading to increased levels of volatility affect all individual stocks identically. There is a unique market volatility factor. Of course, this is a very restrictive assumption since one should expect *a priori* other, sector specific, sources of volatility. This, we believe, is the main reason for the discrepancies between elliptical models and the empirical data reported below.

Some of the above measures of dependence can be explicitly computed for elliptical models. Introducing the following ratios: $f_d = \mathbb{E}[\sigma^{2d}] / \mathbb{E}[\sigma^d]^2$, one readily

finds:

$$\zeta^{(2)}(\rho) = \frac{f_2 \cdot (1 + 2\rho^2) - 1}{3f_2 - 1} \quad (3.2a)$$

$$\zeta^{(1)}(\rho) = \frac{f_1 \cdot D(\rho) - 1}{\frac{\pi}{2}f_1 - 1} \quad (3.2b)$$

$$C^*(\rho) = \frac{1}{4} + \frac{1}{2\pi} \arcsin \rho = \frac{1}{4} (1 + \rho^{(s)}), \quad (3.2c)$$

where $C^* \equiv C(\frac{1}{2}, \frac{1}{2})$. Note that f_2 is related to the kurtosis of the X 's through the relation $\kappa = 3(f_2 - 1)$. The calculation of the tail correlation coefficients depends on the specific form of $P(\sigma)$, for which several choices are possible. We will focus in the following on two of them, corresponding to the Student model and the log-normal model.

3.2. The Student ensemble

When the distribution of the square volatility is inverse Gamma, i.e. $P(u = \sigma^2) \propto e^{-\frac{1}{u}}/u^{1+\nu/2}$, the joint pdf of the returns turns out to have an explicit form [9,14,30]:

$$t_\nu(\mathbf{x}) = \frac{1}{\sqrt{(\nu\pi)^N \det \Sigma}} \frac{\Gamma(\frac{\nu+N}{2})}{\Gamma(\frac{\nu}{2})} \left(1 + \frac{\mathbf{x}^T \Sigma^{-1} \mathbf{x}}{\nu} \right)^{-\frac{\nu+N}{2}} \quad (3.3)$$

This is the multivariate Student distribution with ν degrees of freedom for N random variables with dispersion matrix Σ . Clearly, the marginal distribution of $t_\nu(\mathbf{x})$ is itself a Student distribution with a tail exponent equal to ν , which is well known to describe satisfactorily the univariate pdf of high-frequency returns (from a few minutes to a day or so), with an exponent in the range [3, 5] (see e.g. Refs. [1,8,20,37]). The multivariate Student model is therefore a rather natural choice; its corresponding copula defines the Student copula. For $N = 2$, it is entirely parameterised by ν and the correlation coefficient $\rho = \Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$.

The moments of $P(\sigma)$ are easily computed and lead to the following expressions for the coefficients f_d :

$$f_1 = \frac{2}{\nu - 2} \left(\frac{\Gamma(\frac{\nu}{2})}{\Gamma(\frac{\nu-1}{2})} \right)^2 \quad f_2 = \frac{\nu - 2}{\nu - 4} \quad (3.4a)$$

when $\nu > 2$ (resp. $\nu > 4$). Note that in the limit $\nu \rightarrow \infty$ at fixed N , the multivariate Student distribution boils down to a multivariate Gaussian. The shape of $\zeta^{(1)}(\rho)$ and $\zeta^{(2)}(\rho)$ for $\nu = 5$ is given in Fig. 1.

One can explicitly compute the coefficient of tail dependence for Student variables, which only depends on ν and on the linear correlation coefficient ρ . By symmetry, one has $\tau_{UU}(p; \nu, \rho) = \tau_{LL}(p; \nu, \rho) = \tau_{UL}(p; \nu, -\rho) = \tau_{LU}(p; \nu, -\rho)$. When $p = 1 - \epsilon$ with $\epsilon \rightarrow 0$, one finds:

$$\tau_{UU}(p; \nu, \rho) = \tau_{UU}^*(\nu, \rho) + \beta(\nu, \rho) \cdot \epsilon^{\frac{2}{\nu}} + \mathcal{O}(\epsilon^{\frac{4}{\nu}}), \quad (3.5)$$

where $\tau_{UU}^*(\nu, \rho)$ and $\beta(\nu, \rho)$ are coefficients given in Appendix A, see also Figs. 1 and 2. The notable features of the above results are:

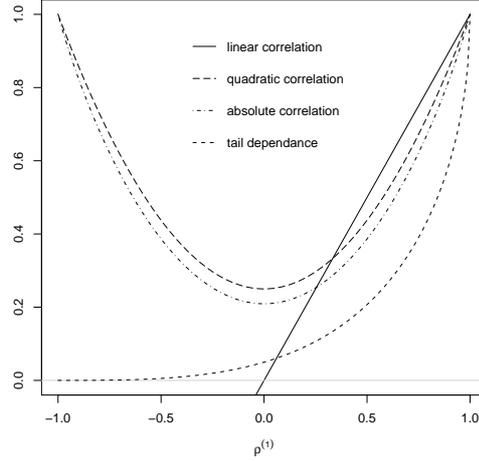


Fig. 1. For all elliptical distributions, the different measures of dependence are functions of ρ . Illustration for a pair with a bivariate Student distribution ($\nu = 5$).

- The asymptotic tail dependence $\tau_{UV}^*(\nu, \rho)$ is strictly positive for all $\rho > -1$ and finite ν , and tends to zero in the Gaussian limit $\nu \rightarrow \infty$ (see Fig. 2(a)). The intuitive interpretation is quite clear: large events are caused by large occasional bursts in volatility. Since the volatility is common to all assets, the fact that one return is positive extremely large is enough to infer that the volatility is itself large. Therefore that there is a non-zero probability $\tau_{UV}^*(\nu, \rho)$ that another asset also has a large positive return (except if $\rho = -1$ since in that case the return can only be large and negative!). It is useful to note that the value of $\tau_{UV}^*(\nu, \rho)$ does not depend on the explicit shape of $P(\sigma)$ provided the asymptotic tail of $P(\sigma)$ decays as $L(\sigma)/\sigma^{1+\nu}$, where $L(\sigma)$ is a slow function.
- The coefficient $\beta(\nu, \rho)$ is also positive, indicating that estimates of $\tau_{UV}^*(\nu, \rho)$ based on measures of $\tau_{UV}(p; \nu, \rho)$ at finite quantiles (e.g. $p = 0.99$) are biased upwards. Note that the correction term is rapidly large because ϵ is raised to the power $2/\nu$. For example, when $\nu = 4$ and $\rho = 0.3$, $\beta \approx 0.263$ and one expects a first-order correction 0.026 for $p = 0.99$. This is illustrated in Figs. 2(a), 2(b). The form of the correction term (in $\epsilon^{2/\nu}$) is again valid as soon as $P(\sigma)$ decays asymptotically as $L(\sigma)/\sigma^{1+\nu}$.
- Not only is the correction large, but the accuracy of the first order expansion is very bad, since the ratio of the neglected term to the first correction is itself $\sim \epsilon^{2/\nu}$. The region where the first term is accurate is therefore exponentially small in ν — see Fig. 2(b).

Finally, we plot in Fig. 3 the rescaled difference between the Student copula

The joint distribution of stock returns is not elliptical 9

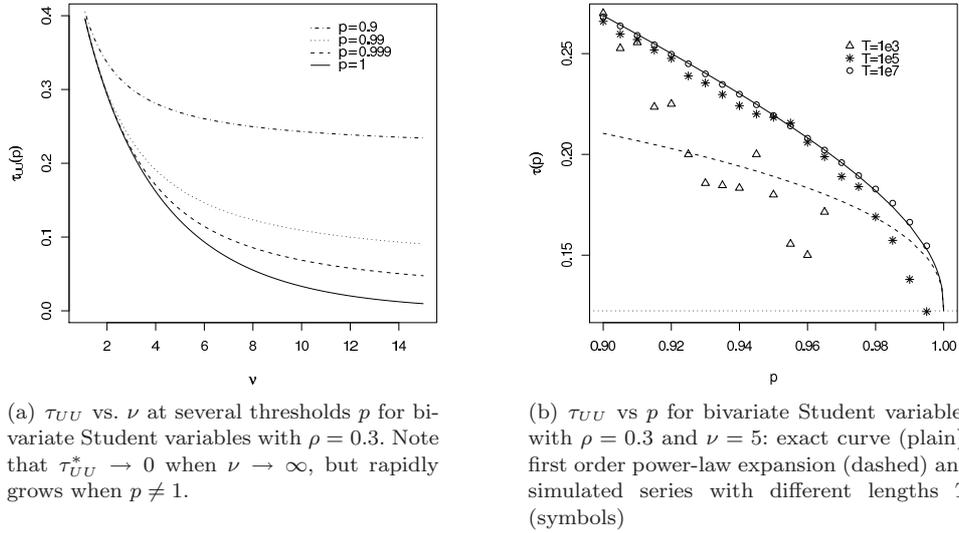


Fig. 2.

and the Gaussian copula both on the diagonal and on the anti-diagonal, for several values of the linear correlation coefficient ρ and for $\nu = 5$. One notices that the difference is zero for $p = \frac{1}{2}$, as expected from the expression of $C^*(\rho)$ for general elliptical models. Away from $p = \frac{1}{2}$ on the diagonal, the rescaled difference has a positive convexity and non-zero limits when $p \rightarrow 0$ and $p \rightarrow 1$, corresponding to τ_{LL}^* and τ_{UU}^* .

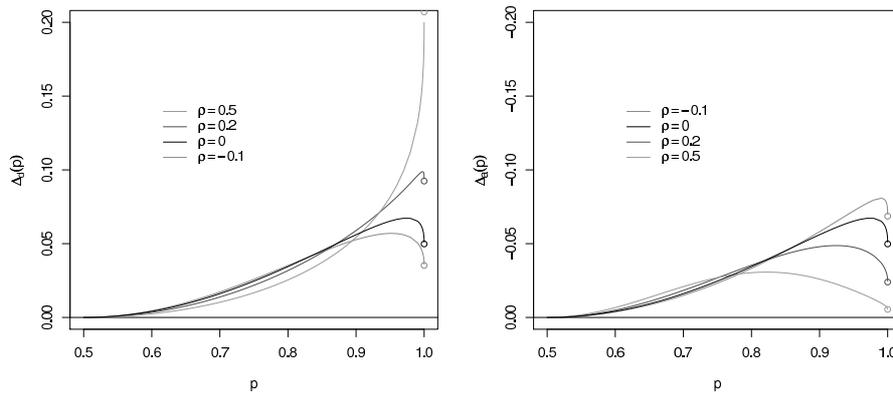


Fig. 3. Diagonal and anti-diagonal of the Student copula: the quantities $\Delta_d(p)$ and $\Delta_a(p)$ defined in Eq. (2.11) are plotted versus p for several values of ρ and fixed $\nu = 5$. (The curves are identical in the range $p \in [0, 0.5]$ due to the symmetry $p \leftrightarrow 1-p$).

3.3. The log-normal ensemble

If we now choose $\sigma = \sigma_0 e^\xi$ with $\xi \sim \mathcal{N}(0, s)$ (as suggested by advocates of multifractal models [3,28,35]), the resulting multivariate model defines the log-normal ensemble and the log-normal copula. The factors f_d are immediately found to be: $f_d = e^{d^2 s^2}$ with no restrictions on d . The Gaussian case now corresponds to $s \equiv 0$.

Although the inverse Gamma and the lognormal distributions are very different, it is well known that the tail region of the log-normal is in practice very hard to distinguish from a power-law. In fact, one may write:

$$\hat{\sigma}^{-1} e^{-\frac{\ln^2 \hat{\sigma}}{2s^2}} = \hat{\sigma}^{-(1+\nu_{\text{eff}})} \quad \nu_{\text{eff}} = \frac{\ln \hat{\sigma}}{2s^2}, \quad (3.6)$$

where we have introduced $\hat{\sigma} = \sigma/\sigma_0$. The above equation means that for large $\hat{\sigma}$ a lognormal is like a power-law with a slowly varying exponent. If the tail region corresponds to $\hat{\sigma} \in [4, 8]$ (say), the effective exponent ν lies in the range $[0.7/s^2, 1.05/s^2]$. Another way to obtain an approximate dictionary between the Student model and the lognormal model is to fix $\nu_{\text{eff}}(s)$ such that the coefficient f_2 (say) of the two models coincide. Numerically, this leads to $\nu_{\text{eff}}(s) \approx 2 + 0.5/s^2$, leading to $s \approx 0.4$ for $\nu = 5$. In any case, our main point here is that from an empirical point of view, one does not expect striking differences between a Student model and a log-normal model — even though from a strict mathematical point of view, the models are very different. In particular, the asymptotic tail dependence coefficients τ_{UU}^* or τ_{LL}^* are always zero for the log-normal model (but $\tau_{LL}(p)$ or $\tau_{UU}(p)$ converge exceedingly slowly towards zero as $p \rightarrow 1$).

3.4. Pseudo-elliptical generalisation

In the previous elliptical description of the returns, all stocks are subject to the exact same stochastic volatility, what leads to non-linear dependences like tail effects and residual higher-order correlations even for $\rho = 0$ (see Fig. 1). In order to be able to fine-tune somewhat this dependence due to the common volatility, a simple generalisation is to let each stock be influenced by an idiosyncratic volatility, thus allowing for a more subtle structure of dependence. More specifically, we write^a:

$$X_i = \sigma_i \cdot \epsilon_i \quad i = 1 \dots N \quad (3.7)$$

^a In vectorial form, we collect the individual (yet dependent) stochastic volatilities σ_i on the diagonal of a matrix D , and $X = D\epsilon$ can be decomposed as $X = RDAU$ where U is a random vector uniformly distributed on the unit hypersphere, the radial component R is a chi-2 random variable independent of U with $\text{rank}(A)$ degrees of freedom, and A is a matrix with appropriate dimensions such that $AA^\dagger = \Sigma$. This description can be contrasted with the one proposed in Ref. [25] under the term ‘‘Multi-tail Generalized Elliptical Distribution’’: $X = R(\vec{u})AU$ with R now depending on the unit vector $\vec{u} = AU/|AU|$. In other words, the description in (3.7) provides a different radial amplitude for each component, whereas [25] characterises a direction-dependent radial part identical for every component. The latter allows for a richer phenomenology than the former, but lacks financial intuition.

where the Gaussian residuals ϵ_i have the same joint distribution as before, and are independent of the σ_i , but we now generalise the definition of the ratios f_d :

$$f_d^{(i,j)} = \frac{\mathbb{E}[\sigma_i^d \sigma_j^d]}{\mathbb{E}[\sigma_i^d] \mathbb{E}[\sigma_j^d]} \quad (3.8)$$

which describe the joint distribution of the volatilities. As an explicit example, we consider the natural generalisation of the log-normal model and write $\sigma_i = \sigma_{0i} e^{\xi_i}$, with^b $\xi_i \stackrel{\text{id}}{\sim} \mathcal{N}(0, s)$, and some correlation structure of the ξ 's: $\mathbb{E}[\xi_i \xi_j] = s^2 c_{ij}$. One then finds:

$$f_d(c) = e^{d^2 s^2 c}. \quad (3.9)$$

Within this setting, the generalisation of coefficients (3.2) can be straightforwardly calculated. Denoting by r the correlation coefficient of ϵ_i and ϵ_j (now different from ρ) we find:

$$\rho(r, c) = \frac{f_1(c)}{f_1(1)} \cdot r \quad (3.10a)$$

$$\zeta^{(2)}(r, c) = \frac{f_2(c)(1 + 2r^2) - 1}{3f_2(1) - 1} \quad (3.10b)$$

$$\zeta^{(1)}(r, c) = \frac{f_1(c)D(r) - 1}{\frac{\pi}{2}f_1(1) - 1} \quad (3.10c)$$

$$C^*(r, c) = \frac{1}{4} + \frac{1}{2\pi} \arcsin r \quad \forall c \quad (3.10d)$$

These formulas are straightforwardly generalised for arbitrary description of the volatilities, using the coefficients f_d defined in (3.8) instead of the explicit expressions given by (3.9). When c is fixed (e.g. for the elliptical case $c = 1$), ρ and r are proportional, and all measures of non-linear dependences can be expressed as a function of ρ . But this ceases to be true as soon as there exists some non trivial structure c in the volatilities. In that case, c and r are “hidden” underlying variables, that can only be reconstructed from the knowledge of ρ , $\zeta^{(2)}$, C^* , assuming of course that the model is accurate.

Notice that the result on $C(\frac{1}{2}, \frac{1}{2})$ is totally independent of the structure of the volatilities^c. Indeed, what is relevant for the copula at the central point is not the amplitude of the returns, but rather the number of $+/-$ events, which is unaffected by any multiplicative scale factor as long as the median of the univariate marginals is nil. An important consequence of this result is that for all elliptical or pseudo-elliptical model, $\rho = 0$ implies that $C^*(\rho = 0) = \frac{1}{4}$. We now turn to empirical data to test the above predictions of elliptical models, in particular this last one.

^bA further generalisation that allows for stock dependent “vol of vol” s_i is also possible.

^cThis property holds even when σ_i depends on the sign of ϵ_i , which might be useful to model the leverage effect that leads to some asymmetry between positive and negative tails, as the data suggests.

3.5. A word on Archimedean copulas

In the universe of all possible copulas, a particular family has become increasingly popular in finance: that of “Archimedean copulas”. These copulas are defined, for bivariate random vectors, as follows [44]:

$$C_A(u, v) \equiv \phi^{-1} [\phi(u) + \phi(v)], \quad (3.11)$$

where $\phi(u) : [0, 1] \rightarrow [0, 1]$ is a function such that $\phi(1) = 0$ and ϕ^{-1} is decreasing and completely monotone. For example, Frank copulas are such that $\phi_F(u) = \ln[e^\theta - 1] - \ln[e^{\theta u} - 1]$ where θ is a real parameter, or Gumbel copulas, such that $\phi_G(u) = (-\ln u)^\theta$, $0 < \theta \leq 1$. The asymptotic coefficient of tail dependence are all zero for Frank copulas, whereas $\tau_{UV}^* = 2 - 2^\theta$ (and all other zero) for the Gumbel copulas. The case of general multivariate copulas is obtained as natural generalisation of the above definition.

One can of course attempt to fit empirical copulas with a specific Archimedean copula. By playing enough with the function ϕ , it is obvious that one will eventually reach a reasonable quality of fit. What we take issue with is the complete lack of intuitive interpretation or plausible mechanisms to justify why the returns of two correlated assets should be described with Archimedean copulas. This is particularly clear after recalling how two Archimedean random variables are generated: first, take two $\mathcal{U}[0, 1]$ random variables s, w . Set $w' = K^{-1}(w)$ with $K(x) = x - \phi(x)/\phi'(x)$. Now, set:

$$u = \phi^{-1} [s\phi(w')]; \quad v = \phi^{-1} [(1-s)\phi(w')], \quad (3.12)$$

and finally write $X_1 = \mathcal{P}_{<,1}^{-1}(u)$ and $X_2 = \mathcal{P}_{<,2}^{-1}(v)$ to obtain the two Archimedean returns with the required marginals \mathcal{P}_1 and \mathcal{P}_2 [21,44]. Unless one finds a natural economic or microstructural interpretation for such a labyrinthine construction, we content that such models should be discarded *a priori*, for lack of plausibility. In the same spirit, one should really wonder why the financial industry has put so much faith in Gaussian copulas models to describe correlation between *positive* default times, that were then used to price CDO's and other credit derivatives. We strongly advocate the need to build models bottom-up: mechanisms should come before any mathematical representation (see also [34] and references therein for a critical view on the use of copulas).

4. Empirical study of the dependence of stock pairs

4.1. Methodology: what we do and what we do not do

As the title stresses, the object of this study is to dismiss the elliptical copula as description of the multivariate dependence structure of stock returns. Concretely, the null-hypothesis “ H_0 : the joint distribution is elliptical” admits as corollary “all pairwise bivariate marginal copulas are elliptical and differ only by their linear correlation coefficient”. In other words, all pairs with the same linear correlation ρ are supposed to have identical values of non-linear dependence measures, and

this value is predicted by the elliptical model. Focusing the empirical study on the pairwise measures of dependence, we will reject H_0 by showing that their average value over all pairs with a given ρ is different from the value predicted by elliptical models.

Our methodology differs from usual hypothesis testing using statistical tools and goodness of fit tests, as can be encountered for example in [29] for testing the Gaussian copula hypothesis on financial assets. Indeed, the results of such tests are often not valid because financial time series are *persistent* (although almost not linearly autocorrelated), so that successive realisations cannot be seen as independent draws of an underlying distribution, see the recent discussion on this issue in [5].

4.2. Data set and time dependence

The dataset we considered is composed of daily returns of 1500 stocks labeled in USD in the period 1995–2009 (15 full years). We have cut the full period in three subperiods (1995–1999, 2000–2004 and 2005–2009), and also the stock universe into three pools (large caps, mid-caps and small caps). We have furthermore extended our analysis to the Japanese stock markets. Qualitatively, our main conclusions are robust and do not depend neither on the period, nor on the capitalisation or the market. Some results, on the other hand, do depend on the time period, but we never found any strong dependence on the capitalisation. All measures of dependence are calculated pairwise with non-parametric estimators, and using all trading dates shared by both equities in the pair.

We first show the evolution of the linear correlation coefficients (Fig. 4) and the upper and lower tail dependence coefficients for $p = 0.95$ (Fig. 5) as a function of time for the large-cap stocks. One notes that (a) the average linear correlation $\bar{\rho}(t)$ fluctuates quite a bit, from around 0.1 in the mid-nineties to around 0.5 during the 2008 crisis; (b) the distribution of correlation coefficients shifts nearly rigidly around the moving average value $\bar{\rho}(t)$; (c) there is a marked, time dependent asymmetry between the average upper and lower tail dependence coefficients; (d) overall, the tail dependence tracks the behaviour of $\bar{\rho}(t)$. More precisely, we show the time behaviour of the tail dependence assuming either a Gaussian underlying copula or a Student ($\nu = 5$) copula with the same average correlation $\bar{\rho}(t)$. Note that the former model works quite well when $\bar{\rho}(t)$ is small, whereas the Student model fares better when $\bar{\rho}(t)$ is large. We will repeatedly come back to this point below.

We computed the quadratic $\zeta^{(2)}$ and absolute $\zeta^{(1)}$ correlation coefficients for each pair in the pool, as well as the whole rescaled diagonal and anti-diagonal of the empirical copulas (this includes the coefficients of tail dependence^d, see Sect. 2.3 above). We then average these observables over all pairs with a given linear coefficient ρ , within bins of varying width in order to take account of the frequency

^dThe empirical relative bias for $p = 0.95$ is $\lesssim 1\%$ even for series with only 10^3 points — typically daily returns over 4 years.

of observations in each bin. We show in Fig. 6 $\zeta^{(1)}$ and $\tau_{UU}(0.95), \tau_{LL}(0.95)$ as a function of ρ for all stocks in the period 2000–2004, together with the prediction of the Student copula model for various values of ν , including the Gaussian case $\nu = \infty$. For both quantities, we see that the empirical curves systematically cross the Student predictions, looking more Gaussian for small ρ 's and compatible with Student $\nu \approx 6$ for large ρ 's, echoing the effect noticed in Fig. 5 above. The same effect would appear if one compared with the log-normal copula: elliptical models imply a residual dependence due to the common volatility, *even when the linear correlation goes to zero*. This property is not observed empirically, since we rather see that higher-order correlations almost vanish together with ρ . The assumption that a common volatility factor affects all stocks is therefore too strong, although it seems to make sense for pairs of stocks with a sufficiently large linear correlation. This result is in fact quite intuitive and we will expand on this idea in Sect. 4.4 below.

The above discrepancies with the predictions of elliptical models are qualitatively the same for all periods, market caps and is also found for the Japan data set. Besides, elliptical models predict a symmetry between upper and lower tails, whereas the data suggest that the tail dependence is asymmetric. Although most of the time the lower tail dependence is stronger, there are periods (such as 2002) when the asymmetry is reversed.

4.3. Copula diagonals

As argued in Sect. 2.3 above, it is convenient to visualise the rescaled difference $\Delta_{d,a}(p)$ between the empirical copula and the Gaussian copula, along the diagonal (p, p) and the anti-diagonal $(p, 1 - p)$, see Eqs. (2.11). The central point $p = \frac{1}{2}$ is special; since $C_G(\frac{1}{2}, \frac{1}{2}) = \frac{1}{4}$, $\Delta_{d,a}(\frac{1}{2})$ are both equal to $4C^* - 1$, that takes a universal value for all elliptical models, given by Eq. (3.2c).

We show in Fig. 8 the diagonal and anti-diagonal copulas for all pairs of stocks in the period 2000–2004, and for various values of ρ . We also show the prediction of the Student $\nu = 5$ model and of a Frank copula model with Student $\nu = 5$ marginals and the appropriate value of ρ . What is very striking in the data is that $\Delta_d(p)$ is concave for small ρ 's and becomes convex for large ρ 's, whereas the Student copula diagonal is *always convex*. This trend is observed for all periods, and all caps, and is qualitatively very similar in Japan as well for all periods between 1991 and 2009. We again find that the Student copula is a reasonable representation of the data only for large enough ρ . The Frank copula is always a very bad approximation – see the wrong curvature along the diagonal and the inaccurate behaviour along the anti-diagonal.

Let us now turn to the central point of the copula, $C^* = C(\frac{1}{2}, \frac{1}{2})$. We plot in Fig. 7 the quantity $-\cos(2\pi C^*) - \rho$ as a function of ρ , which should be zero for all elliptical models, according to Eq. (3.2c). The data here include the 284 equities constantly member of the S&P500 index in the period 2000–2009.

We again find a clear systematic discrepancy, that becomes stronger for smaller ρ 's: the empirical value of C^* is too large compared with the elliptical prediction. In particular, for stocks with zero linear correlation ($\rho = 0$), we find $C^* > \frac{1}{4}$, i.e. even when two stocks are uncorrelated, the probability that they move in the same direction^e is larger than $\frac{1}{4}$. The bias shown in Fig. 7 is again found for all periods and all market caps, and for Japanese stocks as well. Only the amplitude of the effect is seen to change across the different data sets.

Statistical errors in each bin are difficult to estimate since pairs containing the same asset are mechanically correlated. In order to ascertain the significance of the previous finding, we have compared the empirical value of $C^*(\rho)$ with the result of a numerical simulation, where we generate time series of Student ($\nu = 5$) returns using the empirical correlation matrix. In this case, the expected result is that all pairs with equal correlation have the same bivariate copula and thus the same C^* so that the dispersion of the results gives an estimate of measurement noise. We find that, as expected, $C^*(\rho)$ is compatible with Eq. (3.2c), at variance with empirical results. We also find that the dispersion of the empirical points is significantly larger than that of the simulated elliptical pairs with identical linear correlations, suggesting that all pairs *cannot be described by the same bivariate copula* (and definitively not an elliptical copula, as argued above).

All these observations, and in particular the last one, clearly indicate that Student copulas, or any elliptical copulas, are inadequate to represent the full dependence structure in stock markets. Because this class of copulas has a transparent interpretation, we in fact know why this is the case: assuming a common random volatility factor for all stocks is oversimplified. This hypothesis is indeed only plausible for sufficiently correlated stocks, in agreement with the set of observations we made above. As a first step to relax this hypothesis, we now turn to the pseudo-elliptical log-normal model, that allows further insights but still has unrecoverable failures.

4.4. Pseudo-elliptical log-normal model

As we just showed, the fact that the central value of the copula C^* does not obey the relation $-\cos(2\pi C^*) = \rho$ rules out all elliptical models. One possible way out is to consider a model where random volatilities are stock dependent, as explained in Sect. 3.4. Choosing for simplicity a log-normal model of correlated volatilities and inserting (3.9) into (3.10), the new prediction is:

$$\rho = -e^{s^2(c-1)} \cos(2\pi C^*), \quad (4.1)$$

where c is the correlation of the log-volatilities. This suggests to plot $z = \ln(\rho/|\cos(2\pi C^*)|)$ as a function of ρ , as shown in Fig. 7(b). For a purely

^eA more correct statement is that both stocks have returns below their median with probability larger than $\frac{1}{4}$. However, the median of the distributions are very close to zero, justifying our slight abuse of language.

elliptical model, z should be identically zero, corresponding to perfectly correlated volatilities ($c = 1$). What we observe in Fig. 7(b), on the other hand, is that c is indeed close to unity for large enough ρ 's, but systematically decreases as ρ decreases; in other words, the volatilities of weakly correlated stocks are themselves weakly correlated. This is, again, in line with the conclusion we reached above.

However, this pseudo-elliptical model still predicts that $C^* = \frac{1}{4}$ for stocks with zero linear correlations, in disagreement with the data shown in Fig. 7(a). This translates, in Fig. 7(b), into a negative divergence of z when $\rho \rightarrow 0$. This finding means that we should look for other types of constructions. How can one have at the same time $\rho = 0$ and $C^* > \frac{1}{4}$? A toy-model, that serves as the basis for a much richer model that we will report elsewhere [6], is the following. Consider two independent, symmetrically distributed random factors ψ_1 and ψ_2 with equal volatilities, and construct the returns of assets 1 and 2 as:

$$X_{1,2} = \psi_1 \pm \psi_2. \quad (4.2)$$

Clearly, the linear correlation is zero. Using a cumulant expansion, one easily finds that to the kurtosis order, the central value of the copula is equal to:

$$C^* = \frac{1}{4} + \frac{\kappa_2 - \kappa_1}{24\pi} + \dots \quad (4.3)$$

Therefore, if the kurtosis κ_1 of the factor to which both stocks are positively exposed is smaller than the kurtosis κ_2 of the spread, one does indeed find $C^*(\rho = 0) > \frac{1}{4}$.

We will investigate in a subsequent paper [6] additive models such as the above one, with random volatilities affecting the factors ψ_1 and ψ_2 , that generalise elliptical models in a way to capture the presence of several volatility modes.

5. Conclusion

The object of this paper was to discuss the adequacy of Student copulas, or more generally elliptical copulas, to describe the multivariate distribution of stock returns. We have suggested several ideas of methodological interest to efficiently visualise and compare different copulas. We recommend in particular the rescaled difference with the Gaussian copula along the diagonal and the central value of the copula as strongly discriminating observables. We have studied the dependence of these quantities, as well as other non-linear correlation coefficients, with the linear correlation coefficient ρ , and made explicit the predictions of elliptical models that can be empirically tested. We have shown, using a very large data set, with the daily returns of 1500 US stocks over 15 years, that elliptical models fail to capture the detailed structure of stock dependences.

In a nutshell, the main message elicited by our analysis is that Student copulas provide a good approximation to describe the joint distribution of strongly correlated pairs of stocks, but badly miss their target for weakly correlated stocks. We believe that the same results hold for a wider class of assets: it is plausible that highly correlated assets do indeed share the same risk factor. Intuitively, the

failure of elliptical models to describe weakly correlated assets can be traced to the inadequacy of the assumption of a single “market” volatility mode. We expect that exactly as for returns, several factors are needed to capture sectorial volatility modes and idiosyncratic modes as well. The precise way to encode this idea into a workable model that would naturally generalise elliptical models and faithfully capture the non-linear, asymmetric dependence in stock markets is at this stage an open problem, on which we will report in an upcoming paper. We strongly believe that such a quest cannot be based on a formal construction of mathematically convenient models (such as Archimedean copulas that, in our opinion, cannot be relevant to describe asset returns). The way forward is to rely on intuition and plausibility and come up with models that make financial sense. This is, of course, not restricted to copula modeling, but applies to all quarters of quantitative finance.

Appendix A. Proof of Eqn. (3.5), page 7

Lemma Appendix A.1. *Let (X_1, X_2) follow a bivariate student distribution with ν degrees of freedom and correlation ρ . Denote by $T_\nu(x)$ the univariate Student cdf with ν degrees of freedom, and $x_p \equiv T_\nu^{-1}(p)$ its p -quantile, and define*

$$K = \sqrt{\frac{\nu+1}{1-\rho^2}} \frac{X_1 - \rho x_p}{\sqrt{\nu + X_2^2}} \quad (\text{A.1a})$$

$$k(p) = \frac{\sqrt{\nu+1}\sqrt{1-\rho}}{\sqrt{1+\rho}} (\nu \cdot x_p^{-2} + 1)^{-\frac{1}{2}} = \frac{k(1)}{\sqrt{1 + \frac{\nu}{x_p^2}}} \quad (\text{A.1b})$$

Then,

$$\mathbb{P}[K \leq k(p) \mid X_2 = x_p] = T_{\nu+1}(k(p))$$

Proof. The proof proceeds straightforwardly by showing that $t_{\nu+1}(k) = \mathcal{P}_{X_1|X_2}(x_1|x_2) \frac{\partial x_1}{\partial k}$ when $x_2 = x_p$. The particular result for the limit case $p = 1$ is stated in [14]. \square

Theorem Appendix A.1. ^f *Let (X_1, X_2) be defined like in the Lemma above. The pre-asymptotic behaviour of its tail dependence when $p \rightarrow 1$ is approximated by the following expansion in (rational) powers of $(1-p)$:*

$$\tau_{UV}(p) = \tau_{UV}^*(\nu, \rho) + \beta(\nu, \rho) \cdot (1-p)^{\frac{2}{\nu}} + \mathcal{O}((1-p)^{\frac{4}{\nu}}) \quad (3.5)$$

^fThis result was simultaneously found by Manner and Segers [31] (in a somewhat more general context). We still sketch our proof because it follows a different route (uses the Copula), and the final expression looks quite different, although of course numerically identical.

18 *R. Chicheportiche and J.-P. Bouchaud*

with

$$\tau_{UV}^*(\nu, \rho) = 2 - 2 T_{\nu+1}(k(1)) \quad (\text{A.2})$$

$$\begin{aligned} \beta(\nu, \rho) &= \frac{\nu^{\frac{2}{\nu}+1}}{\frac{2}{\nu}+1} k(1) t_{\nu+1}(k(1)) L_\nu^{-\frac{2}{\nu}} \quad (\text{A.3}) \\ &= \left(\frac{\Gamma(\frac{\nu}{2})\sqrt{\pi}}{\Gamma(\frac{\nu+1}{2})} \right)^{\frac{2}{\nu}} \frac{\nu^{\frac{2}{\nu}}}{\frac{2}{\nu}+1} \frac{\sqrt{\nu+1}\sqrt{1-\rho}}{\sqrt{1+\rho}} \cdot t_{\nu+1} \left(\frac{\sqrt{\nu+1}\sqrt{1-\rho}}{\sqrt{1+\rho}} \right) \end{aligned}$$

Proof.

Recall from Eq. (2.9) that

$$\tau_{UV}(p) = 2 - \frac{C(p, p) - C(1, 1)}{p - 1} = 2 - \frac{1}{1 - p} \int_p^1 \frac{dC(p, p)}{dp} dp \quad (\text{A.4})$$

One easily shows that

$$dC(p, p) = 2 \mathbb{P}[X_1 \leq x_p \mid X_2 = x_p] dp = 2 T_{\nu+1}(k(p)) dp$$

where the second equality holds in virtue of the aforementioned lemma.

Now, for p close to 1, $k(p) = k(1) \left(1 - \frac{1}{2} \frac{\nu}{x_p^2}\right) + \mathcal{O}(x_p^{-4})$, and

$$\frac{dC(p, p)}{dp} = 2 T_{\nu+1}(k(p)) \approx 2 T_{\nu+1}(k(1)) - 2(k(1) - k(p)) \cdot t_{\nu+1}(k(1))$$

But since the Student distribution behaves as a power-law precisely in the region $p \lesssim 1$, we write $t_\nu(x) \approx L_\nu/x^{\nu+1}$ and immediately get

$$1 - p = \frac{L_\nu}{\nu} x_p^{-\nu} \quad \text{with} \quad L_\nu = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi}} \nu^{\frac{\nu}{2}}$$

The result follows by collecting all the terms and performing the integration in (A.4). \square

Notice that for large ν , the exponent is almost zero and the correction term is of order $\mathcal{O}(1)$, and the expansion ceases to hold. This is particularly true for the Gaussian distribution ($\nu = \infty$) for which the behaviour is radically different at the limit $p \rightarrow 1$ (where there's strictly no tail correlation) and at $p < 1$ where a dependence subsists, see Fig. 2(a).

Acknowledgments

We want to thank F. Abergel, R. Allez, S. Ciliberti, L. Laloux & M. Potters for help and useful comments.

References

- [1] J.-P. Bouchaud and M. Potters, *Theory of Financial Risk and Derivative Pricing: from Statistical Physics to Risk Management*, Cambridge University Press (2003).
- [2] D. Brigo, A. Pallavicini and R. Torresetti, *Credit Models and the Crisis: A journey into CDOs, Copulas, Correlations and Dynamic Models*, Wiley, Chichester (2010).
- [3] L. E. Calvet and A. Fisher, *Multifractal Volatility: Theory, Forecasting, and Pricing*, Academic Press (2008).
- [4] S. Cambanis, S. Huang and G. Simons, On the theory of elliptically contoured distributions, *Journal of Multivariate Analysis* **11** (1981) 368–385.
- [5] R. Chicheportiche and J.-P. Bouchaud, Goodness-of-fit tests with dependent observations, *Journal of Statistical Mechanics: Theory and Experiment* **2011** (2011) P09003.
- [6] R. Chicheportiche and J.-P. Bouchaud, In preparation, *In preparation* (2011).
- [7] S. Coles, J. Heffernan and J. Tawn, Dependence Measures for Extreme Value Analyses, *Extremes* **2** (1999) 339–365.
- [8] R. Cont, Empirical properties of asset returns: stylized facts and statistical issues, *Quantitative Finance* **1** (2001) 223–236.
- [9] S. Demarta and A. J. McNeil, The t copula and related copulas, *International Statistical Review* **73** (2005) 111–129.
- [10] M. A. H. Dempster (ed.), *Risk Management: Value at Risk and Beyond*, Cambridge University Press, Cambridge (2002).
- [11] B. Dupire, Pricing with a smile, *Risk* **7** (1994) 18–20.
- [12] V. Durrleman, A. Nikeghbali and T. Roncalli, Which copula is the right one, *SSRN working paper* (2000).
- [13] N. El Karoui, High-dimensionality effects in the Markowitz problem and other quadratic programs with linear equality constraints: risk underestimation, *Working paper*, University of Berkeley (2009).
- [14] P. Embrechts, A. J. McNeil and D. Straumann, *Correlation And Dependence In Risk Management: Properties And Pitfalls*, chapter 7, In Dempster [10] (2002) 176–223.
- [15] J.-D. Fermanian and O. Scaillet, *Some Statistical Pitfalls in Copula Modeling for Financial Applications*, chapter 4, In Klein [24] (2005) 57–72.
- [16] M. Fischer, C. Köck, S. Schlüter and F. Weigert, An empirical analysis of multivariate copula models, *Quantitative Finance* **9** (2009) 839–854.
- [17] I. Fortin and C. Kuzmics, Tail-dependence in stock-return pairs, *International Journal of Intelligent Systems in Accounting, Finance & Management* **11** (2002) 89–107.
- [18] G. Frahm, M. Junker and A. Szimayer, Elliptical copulas: applicability and limitations, *Statistics & Probability Letters* **63** (2003) 275–286.
- [19] R. Frey, A. J. McNeil and M. A. Nyfeler, Modelling Dependent Defaults: Asset Correlations Are Not Enough!, *Technical report*, Swiss Banking Institute, University of Zurich, ETHZ (March 2001).
- [20] M. A. Fuentes, A. Gerig and J. Vicente, Universal Behavior of Extreme Price Movements in Stock Markets, *Public Library of Science ONE* **4** (2009) e8243.
- [21] C. Genest and L.-P. Rivest, Statistical inference procedures for bivariate Archimedean copulas, *Journal of the American Statistical Association* **88** (1993) 1034–1043.
- [22] P. S. Hagan, D. Kumar, A. S. Lesniewski and D. E. Woodward, Managing Smile Risk, *Wilmott Magazine* (2002) 84–108.
- [23] H. Hult and F. Lindskog, Multivariate extremes, aggregation and dependence in elliptical distributions, *Advances in Applied probability* **34** (2002) 587–608.
- [24] E. Klein (ed.), *Capital formation, governance and banking*, Nova Publishers (2005).
- [25] S. Kring, S. T. Rachev, M. Höchstötter, F. J. Fabozzi and M. L. Bianchi, Multi-tail generalized elliptical distributions for asset returns, *Econometrics Journal* **12** (2009)

- 272–291.
- [26] O. Ledoit and M. Wolf, A well-conditioned estimator for large-dimensional covariance matrices, *Journal of multivariate analysis* **88** (2004) 365–411.
 - [27] X. Luo and P. V. Shevchenko, The t copula with multiple parameters of degrees of freedom: bivariate characteristics and application to risk management, *Quantitative Finance* **9** (2009) 1–16.
 - [28] T. Lux, The Multi-Fractal Model of Asset Returns: Its Estimation via GMM and Its Use for Volatility Forecasting, *Journal of Business and Economic Statistics* **26** (2008) 194.
 - [29] Y. Malevergne and D. Sornette, Testing the Gaussian copula hypothesis for financial assets dependences, *Quantitative Finance* **3** (2003) 231–250.
 - [30] Y. Malevergne and D. Sornette, *Extreme financial risks: From dependence to risk management*, Springer Verlag (2006).
 - [31] H. Manner and J. Segers, Tails of correlation mixtures of elliptical copulas, *Insurance: Mathematics and Economics* .
 - [32] M. Marsili, Dissecting financial markets: sectors and states, *Quantitative Finance* **2** (2002) 297.
 - [33] R. Mashal and A. Zeevi, Beyond correlation: Extreme co-movements between financial assets, *SSRN working paper* (2002).
 - [34] T. Mikosch, Copulas: Tales and facts, *Extremes* **9** (2006) 3–20.
 - [35] J.-F. Muzy, E. Bacry and J. Delour, Modelling fluctuations of financial time series: from cascade process to stochastic volatility model, *The European Physical Journal B* **17** (2000) 537–548.
 - [36] A. Patton, Estimation of Copula Models for Time Series of Possibly Different Length, *Economics working paper series*, University of California at San Diego (2001).
 - [37] V. Plerou, P. Gopikrishnan, L. A. N. Amaral, M. Meyer and H. E. Stanley, Scaling of the distribution of price fluctuations of individual companies, *Physical Review E* **60** (1999) 6519.
 - [38] M. Potters and J.-P. Bouchaud, Financial applications of random matrix theory: a short review, *Arxiv preprint q-fin.ST/0190.1205* .
 - [39] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press (2007).
 - [40] W. T. Shaw and K. T. A. Lee, Copula Methods vs Canonical Multivariate Distributions: the multivariate Student T Distribution with general degrees of freedom, *Working paper*, King’s College London (2007).
 - [41] A. Sklar, Fonctions de répartition à n dimensions et leurs marges, *Publ. Inst. Statist. Univ. Paris* **8** (1959) 229–231.
 - [42] V. Tola, F. Lillo, M. Gallegati and R. N. Mantegna, Cluster analysis for portfolio optimization, *Journal of Economic Dynamics and Control* **32** (2008) 235–258.
 - [43] M. Tumminello, F. Lillo and R. N. Mantegna, Hierarchically nested factor model from multivariate data, *EPL (Europhysics Letters)* **78** (2007) 30006.
 - [44] F. Wu, E. A. Valdez and M. Sherris, Simulating Exchangeable Multivariate Archimedean Copulas and its Applications, *Communications in Statistics - Simulation and Computation* **36** (2007) 1019–1034.

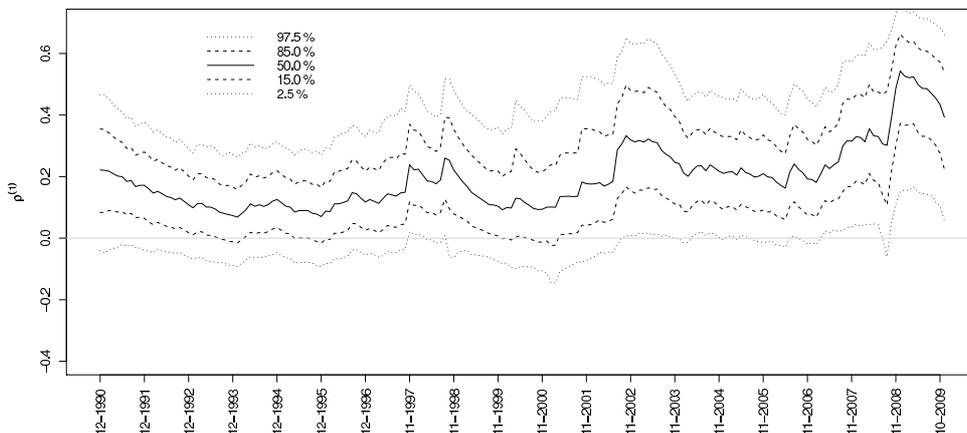


Fig. 4. Time evolution of the coefficient of linear correlation, computed for the stocks in the S&P500 index, with an exponentially moving average of 125 days from January 1990 to December 2009. Five quantiles of the ρ distribution are plotted according to the legend, showing that this distribution moves quasi-rigidly around its median value.

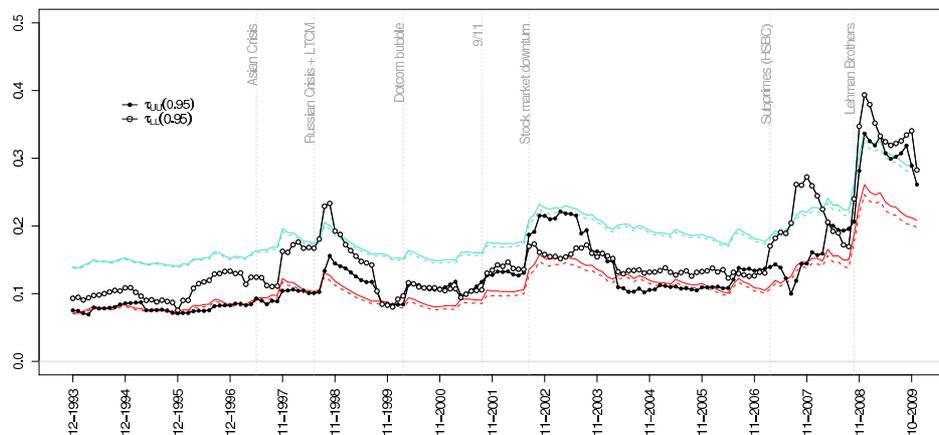


Fig. 5. Time evolution of the coefficient of tail dependence at $p = 0.95$, computed for the stocks in the S&P500 index. A sliding flat window of 250 days moves monthly (25 days) from January 1990 to December 2009. We depict in turquoise and red the predictions for Student pair with $\nu = 5$ and $\nu = \infty$ respectively, based on the evolution of the mean value $\bar{\rho}(t)$ of the cross-sectional linear correlation (dashed) or averaged over the full instantaneous distribution of $\rho(t)$ (plain), with little difference.

22 *R. Chicheportiche and J.-P. Bouchaud*

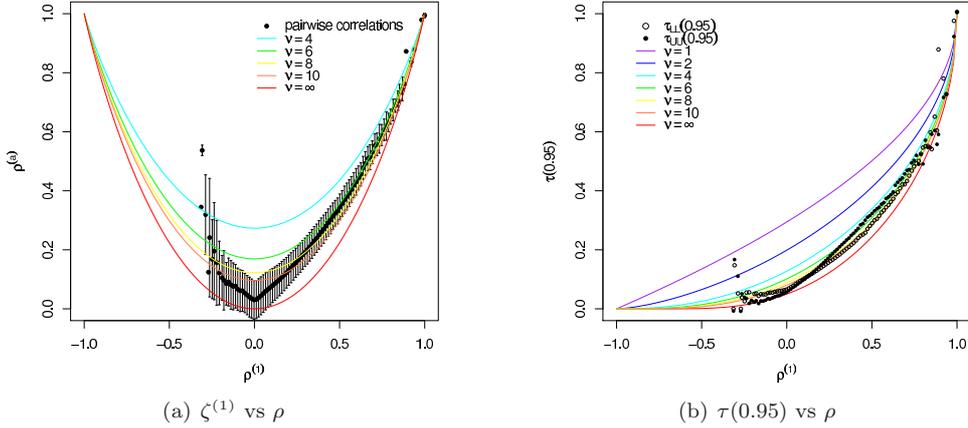


Fig. 6. Empirical (2000–2004) and elliptical. The fact that the empirical points “cross” the theoretical lines corresponding to elliptical models reveals that there is a dependence structure in the volatility, more complex than what is suggested by a single common stochastic scale factor.

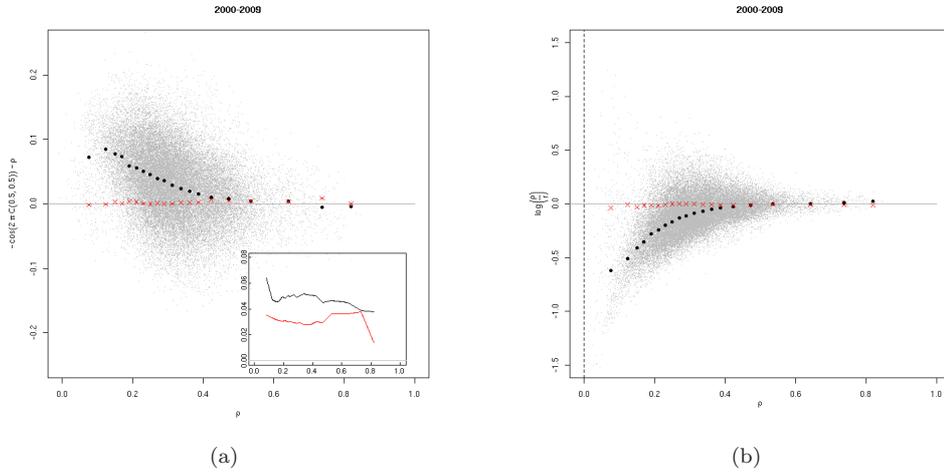


Fig. 7. **Left:** $-\cos(2\pi C^*) + \rho$ versus the coefficient ρ of linear correlation. The grey cloud is a scatter plot of all pairs of stocks. Black dots correspond to the average value of empirical measurements within a correlation bin, and red crosses correspond to the numerically generated elliptical (Student, $\nu = 5$) data, compatible with the expected value zero. Statistical error are not shown, but their amplitudes is of the order of the fluctuations of the red crosses, so that empirics and elliptical prediction don’t overlap for $\rho \lesssim 0.4$. This representation allows to visualise very clearly the systematic discrepancies for small values of ρ , beyond the average value. Inset: the 1.s.d. dispersion of the scattered points inside the bins. **Right:** $z = \ln(\rho/|\cos(2\pi C^*)|)$ as a function of ρ , with the same data and the same conventions. This shows more directly that the volatility correlations of weakly correlated stocks is overestimated by elliptical models, whereas strongly correlated stocks are compatible with the elliptical assumption of a common volatility factor.

The joint distribution of stock returns is not elliptical 23

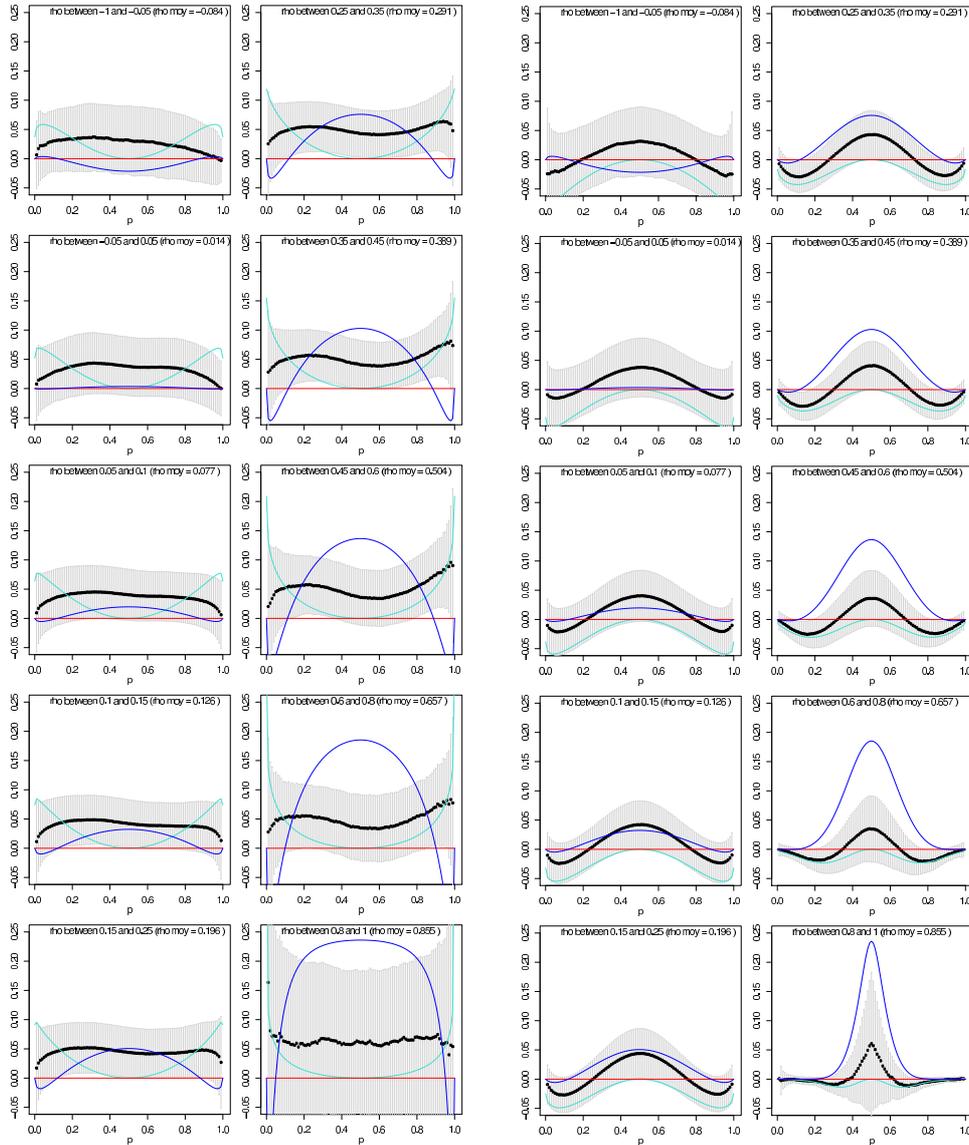


Fig. 8. Diagonal (left) and anti-diagonal (right) of the copula vs p . Real data (2000-2004) averaged in 10 bins of p . The vertical lines corresponds to 1-sigma dispersion of the results (i.e. not the statistical error bar). The turquoise line corresponds to an elliptical model (Student with $\nu = 5$): it predicts too low a value at the centre point and too large tail dependences (left and right limits). The blue line depicts the behaviour of Frank's copula. It is worth noticing that the tail dependence coefficient is always measured empirically at some $p < 1$, for example $p = 0.95$. In this case, the disagreement between data and an elliptical model may be accidentally small. It is important to distinguish the interpretation of the limit coefficient from that of a penultimate value [6].