Goodness-of-fit tests with dependent observations

Rémy Chicheportiche[1,2] **and Jean-Philippe Bouchaud**[1]

[1] Capital Fund Management, 6–8 boulevard Haussmann, 75 009 Paris, France
[2] Chair of Quantitative Finance, Laboratory of Applied Mathematics and Systems, Ecole Centrale Paris, 92 290 Châtenay-Malabry, France

E-mail: remy.chicheportiche@ecp.fr, jean-philippe.bouchaud@cfm.fr

Abstract. We revisit the Kolmogorov-Smirnov and Cramér-von Mises goodness-of-fit (GoF) tests and propose a generalisation to identically distributed, but dependent univariate random variables. We show that the dependence leads to a reduction of the "effective" number of independent observations. The generalised GoF tests are not distribution-free but rather depend on all the lagged bivariate copulas. These objects, that we call "self-copulas", encode all the non-linear temporal dependences. We introduce a specific, log-normal model for these self-copulas, for which a number of analytical results are derived. An application to financial time series is provided. As is well known, the dependence is to be long-ranged in this case, a finding that we confirm using self-copulas. As a consequence, the acceptance rates for GoF tests are substantially higher than if the returns were iid random variables.

Keywords: Extreme value statistics, Stochastic processes, Models of financial markets

PACS numbers: 89.65.Gh, 02.50, 05.45.Tp

AMS classification scheme numbers: 91B84, 62P20, 62M10, 60F05

1. Introduction

Goodness-of-Fit (GoF) tests are designed to assess quantitatively whether a sample of N observations can statistically be seen as a collection of N *independent* realizations of a given probability law, or whether two such samples are drawn from the same hypothetical distribution. Two well-known and broadly used tests in this class are the Kolmogorov-Smirnov (KS) and Cramér-von Mises (CM) tests, that both quantify how close an empirical cumulative distribution function (cdf) F_N is from a target theoretical cdf F (or from another empirical cdf) — see Ref. [1] for a nice review and many references. The major strength of these tests lies in the fact that the asymptotic distributions of their test statistics is completely independent of the null-hypothesis cdf.

It however so happens that in certain fields (physics, finance, geology, etc.) the random variable under scrutiny has some memory. Whereas the unconditional law of the variable may well be unique and independent of time, the conditional probability distribution of an observation *following* a previous observation exhibits particular patterns, and in particular long-memory, even when the linear correlation is short-ranged or trivial. Examples of such phenomena can be encountered in fluid mechanics (the velocity of a turbulent fluid) and finance (stock returns have small auto-correlations but exhibit strong volatility clustering, a form of heteroscedasticity). The long-memory nature of the underlying processes makes it inappropriate to use standard GoF tests in these cases. Still, the determination of the unconditional distribution of returns is a classic problem in quantitative finance, with obvious applications to risk control, portfolio optimization or derivative pricing. Correspondingly, the distribution of stock returns (in particular the behaviour of its tails) has been the subject of numerous empirical and theoretical papers (see e.g. [2, 3] and for reviews [4, 5] and references therein). Clearly, precise statements are only possible if meaningful GoF tests are available.

As a tool to study the — possibly highly non-linear — correlations between returns, "copulas" have long been used in actuarial sciences and finance to describe and model cross-dependences of assets, often in a risk management perspective [6, 7, 5]. Although the widespread use of simple analytical copulas to model multivariate dependences is more and more criticized [8, 9], copulas remain useful as a tool to investigate empirical properties of multivariate data [9].

More recently, copulas have also been studied in the context of auto-dependent univariate time series, where they find yet another application range: just as Pearson's ρ coefficient is commonly used to measure both linear cross-dependences and temporal correlations, copulas are well-designed to assess non-linear dependences both across assets or in time [10, 11, 12] — we will speak of "self-copulas" in the latter case. Interestingly, when trying to extend GoF tests to dependent variables, self-copulas appear naturally. In our empirical study of financial self-copulas, we rely on a non-parametric estimation rather than imposing, for example, a Markovian structure of the

underlying process, as in e.g. [13, 11].

The organisation of the paper is in three parts. In Section 2 we study theoretically how to account for general dependence in GoF tests: we first describe the statistical properties of the empirical cdf of a non-iid vector of observations of finite size, as well as measures of its difference with an hypothesized cdf. We then study the limit properties of this difference and the asymptotic distributions of two norms. In Section 3 we go through a detailed example when the dependences are weak and described by a pseudo-elliptical copula. Section 4 is dedicated to an application of the theory to the case of financial data: after defining our data set, we perform an empirical study of dependences in series of stock returns, and interpret the results in terms of the 'self-copula'; implication of the dependences on GoF tests are illustrated for this special case using Monte-Carlo simulations. The concluding section summarizes the main ideas of the paper, and technical calculations of sections 2 and 3 are collected in the appendix.

2. Goodness-of-fit tests for a sample of dependant drawings

2.1. Empirical cumulative distribution and its fluctuations

Let X be a latent random vector with N identically distributed but dependant variables, with marginal cdf F. One realization of X consists of a time series $\{x_1, \ldots, x_n, \ldots, x_N\}$ that exhibits some sort of persistence. For a given number x in the support of F, let $Y(x)$ be the random vector the components of which are the Bernoulli variables $Y_n(x) = \mathbb{1}_{\{X_n \leq x\}}$. The expectation value and the covariance of $Y_n(x)$ are given by:

$$\mathbb{E}[Y_n(x)] = F(x), \tag{1}$$

$$\mathrm{Cov}(Y_n(x), Y_m(x')) = F_{nm}(x, x') = C_{nm}(F(x), F(x')), \tag{2}$$

where by definition C_{nm} is the "copula" of the random pair (X_n, X_m). The centered mean of $Y(x)$ is:

$$\overline{Y}(x) = \frac{1}{N}\sum_{n=1}^{N} Y_n(x) - F(x) = \langle Y_n(x)\rangle_n - F(x) \tag{3}$$

which measures the difference between the empirically determined cumulative distribution function at point x and its true value. It is therefore the quantity on which any statistics for Goodness-of-Fit testing is built. Denoting $u = F(x), v = F(x')$, the covariance function of $\overline{Y}$ is easily shown to be:

$$\mathrm{Cov}(\overline{Y}(u), \overline{Y}(v)) = \frac{1}{N}(\min(u, v) - uv)\,[1 + \Psi_N(u, v)] \tag{4}$$

where

$$\Psi_N(u, v) = \frac{1}{N}\sum_{n,m\neq n}^{N} \frac{C_{nm}(u, v) - uv}{\min(u, v) - uv} \tag{5}$$

measures the departure from the independent case, corresponding to $C_{nm}(u, v) = uv$ (in which case $\Psi_N(u, v) \equiv 0$). Note that decorrelated but dependant variables may lead to a non zero value of Ψ_N, since the whole pairwise copula enters the formula and not only

the linear correlation coefficients. When the denominator is zero, the fraction should be understood in a limit sense; we recall in particular that [9]

$$\Delta_{nm}(u,u) \equiv \frac{C_{nm}(u,u)-u^2}{u(1-u)} = \tau_{nm}^{UU}(u) + \tau_{nm}^{LL}(1-u) - 1 \tag{6}$$

tends to the upper/lower tail dependence coefficients $\tau_{nm}^{UU}(1)$ and $\tau_{nm}^{LL}(1)$ when $u \to 1$ resp. 0. Intuitively, the presence of $\Psi_N(u,v)$ in the covariance of $\overline{Y}$ above leads to a reduction of the number of effectively independent variables, but a more precise statement requires some further assumptions that we detail below.

In the following, we will restrict to the case of stationary random vectors, for the copula C_{nm} only depends on the lag $t = m - n$. The average of Δ_{nm} over n, m can be turned into an average over t:

$$\Psi_N(u,v) = \sum_{t=1}^{N-1} (1 - \frac{t}{N})(\Delta_t(u,v) + \Delta_{-t}(u,v)) \tag{7}$$

with $\Delta_t(u,v) = \Delta_{n,n+t}(u,v)$. Note that in general $\Delta_t(u,v) \neq \Delta_{-t}(u,v)$, but clearly $\Delta_t(u,v) = \Delta_{-t}(v,u)$, which implies that $\Psi_N(u,v)$ is symmetric in $u \leftrightarrow v$.

We will assume in the following that the dependence encoded by $\Delta_t(u,v)$ has a limited range in time, or at least that it decreases sufficiently fast for the above sum to converge when $N \to \infty$. If the characteristic time scale for this dependence is T, we assume in effect that $T \ll N$. In the example worked out in Section 3 below, one finds:

$$\Delta_t(u,v) = f\left(\frac{t}{T}\right) \frac{A(u,v)}{I(u,v)}, \qquad I(u,v) \equiv \min(u,v) - uv$$

where $f(\cdot)$ is a certain function. If $f(r)$ decays faster than r^{-1}, one finds (in the limit $T \gg 1$):

$$\Psi_\infty(u,v) = \lim_{N\to\infty} \Psi_N(u,v) = T\, \frac{A(u,v)+A(v,u)}{I(u,v)} \int_0^\infty \mathrm{d}r f(r),$$

with corrections at least of the order of T/N when $N \gg T$.

2.2. Limit properties

We now define the process $\tilde{y}(u)$ as the limit of $\sqrt{N}\,\overline{Y}(u)$ when $N \to \infty$. For a given u, it represents the asymptotics of the difference between the empirically determined cdf of the underlying X's and the theoretical one, at the u-th quantile. According to Central Limit Theorem under weak dependences, it is Gaussian as long as the strong mixing coefficients,

$$\alpha_{SM}(t) = \sup_\tau \sup_{A,B} \{|\mathbb{P}(A\cap B) - \mathbb{P}(A)\mathbb{P}(B)| : A \in \sigma(\{Z_n(u)\}_{n\leq\tau}), B \in \sigma(\{Z_n(v)\}_{n\geq\tau+t})\}$$

associated to the sequence $\{Z_n(u)\} = \{Y_n(u) - u\}$, vanish at least as fast as $\mathcal{O}(t^{-5})$‡. We will assume this condition to hold in the following. For example, this condition is met if the function $f(r)$ defined above decays exponentially, or if $f(r \geq 1) = 0$.

The covariance of the process $\tilde{y}(u)$ is given by:

$$H(u,v) = \lim_{N\to\infty} N \operatorname{Cov}(\overline{Y}(u), \overline{Y}(v)) = (\min(u,v) - uv)\,[1 + \Psi_\infty(u,v)] \quad (8)$$

and characterizes a Gaussian bridge since $\mathbb{V}[\tilde{y}(0)] = \mathbb{V}[\tilde{y}(1)] = 0$, or equivalently $\mathbb{P}[\tilde{y}(0) = y] = \mathbb{P}[\tilde{y}(1) = y] = \delta(y)$. Indeed, $I(u,v) = \min(u,v) - uv$ is the covariance function of the Brownian bridge, and $\Psi_\infty(u,v)$ is a non-constant scaling term.

By Mercer's theorem, the covariance $H(u,v)$ can be decomposed on its eigenvectors and $\tilde{y}(u)$ can correspondingly be written as an infinite sum of Gaussian variables:

$$\tilde{y}(u) = \sum_{j=1}^{\infty} U_j(u)\sqrt{\lambda_j}\, z_j \quad (9)$$

where z_j are independent centered, unit-variance Gaussian variables, and the functions U_j and the numbers λ_j are solutions to the eigenvalue problem:

$$\int_0^1 H(u,v)U_i(v)\,\mathrm{d}v = \lambda_i\, U_i(u) \quad \text{with} \quad \int_0^1 U_i(u)U_j(u)\,\mathrm{d}u = \delta_{ij}. \quad (10)$$

In order to measure a limit distance between distributions, a norm over the space of continuous bridges needs to be chosen. Typical such norms are the norm-2 (sum of squares, as the bridge is always integrable), and the norm-sup (as the bridge always reaches an extremal value). In practice, for every given problem, the covariance function in Equation (8) has a specific shape, since $\Psi_\infty(u,v)$ is copula-dependent. Therefore, contrarily to the case of independent random variables, the GoF tests will not be characterized by universal (problem independent) distributions.

2.3. Law of the norm-2 (Cramér-von-Mises)

The norm-2 of the limit process is the integral of $\tilde{y}^2$ over the whole domain:

$$CM = \int_0^1 \tilde{y}(u)^2\,\mathrm{d}u. \quad (11a)$$

In the representation (9), it has a simple expression:

$$CM = \sum_{j=1}^{\infty} \lambda_j z_j^2. \quad (11b)$$

and its law is thus the law of an infinite sum of squared independent gaussian variables weighted by the eigenvalues of the covariance function. Diagonalizing H is thus sufficient

‡ This condition means that the occurence of any two realizations of the underlying variable can be seen as independent for sufficiently long time between the realizations. Since the copula induces a measure of probability on the Borel sets, it amounts in essence to checking that $|C_t(u,v) - uv|$ converges quickly towards 0. See Refs. [14, 15, 10] for definitions of $\alpha-$, $\beta-$, $\rho-$mixing coefficients and sufficient conditions on copulas for geometric mixing (fast exponential decay) in the context of copula-based stationary Markov chains.

to find the distribution of CM, in the form of the Fourier transform of the characteristic function

$$\phi(t) = \mathbb{E}\left[\mathrm{e}^{\mathrm{i}t\,CM}\right] = \prod_j (1 - 2\mathrm{i}t\lambda_j)^{-\frac{1}{2}}\,. \tag{12}$$

The hard task consists in finding the infinite spectrum of H (or some approximations, if necessary).

Ordering the eigenvalues by decreasing amplitude, Equation (11b) makes explicit the decomposition of CM over contributions of decreasing importance so that, at a wanted level of precision, only the most relevant terms can be retained. In particular, if the top eigenvalue dominates all the others, we get the chi-square law with a single degree of freedom:

$$\mathbb{P}[CM \leq k] = \mathrm{erf}\sqrt{\frac{k}{\lambda_0}}. \tag{13}$$

Even if the spectrum cannot easily be determined but $H(u, v)$ is known, all the moments of the distribution can be computed exactly. For example:

$$\mathbb{E}[CM] = \mathrm{Tr}H = \int_0^1 H(u, u)\,\mathrm{d}u, \tag{14a}$$

$$\mathbb{V}[CM] \equiv 2\,\mathrm{Tr}H^2 = 2\int\!\!\int_0^1 H(u, v)^2\,\mathrm{d}u\,\mathrm{d}v. \tag{14b}$$

2.4. Law of the supremum (Kolmogorov-Smirnov)

The supremum of the difference between the empirical cdf of the sample and the target cdf under the null-hypothesis has been used originally by Kolmogorov and Smirnov as the measure of distance. The variable

$$KS = \sup_{u\in[0,1]} |\tilde{y}(u)| \tag{15}$$

describes the limit behaviour of the GoF statistics. In the case where $1 + \Psi_\infty(u, v)$ can be factorized as $\sqrt{\psi(u)}\sqrt{\psi(v)}$, the procedure for obtaining the limiting distribution was worked out in [16], and leads to a problem of a diffusive particle in an expanding cage, for which some results are known. There is however no general method to obtain the distribution of KS for an arbitrary covariance function H.

Nevertheless, if H has a dominant mode, the relation (9) becomes approximately: $\tilde{y}(u) = U_0(u)\sqrt{\lambda_0}z_0 \equiv \kappa_0(u_0)z_0$, and

$$KS = \sqrt{\lambda_0}|z_0| \sup_{u\in[0,1]} |U_0(u)| \equiv \kappa_0(u_0^*)|z_0|. \tag{16}$$

The cumulative distribution function is then simply

$$\mathbb{P}[KS \leq k] = \mathrm{erf}\left(\frac{k}{\sqrt{2}\kappa_0(u_0^*)}\right), \qquad k \geq 0. \tag{17}$$

This approximation is however not expected to work for small values of k, since in this case z_0 must be small, and the subsequent modes are not negligible compared to the

first one. A perturbative correction — working also for large k only — can be found when the second eigenvalue is small, or more precisely when $\tilde{y}(u) = \kappa_0(u)z_0 + \kappa(u)z_1$ with $\epsilon = \kappa/\kappa_0 \ll 1$. The first thing to do is find the new supremum

$$u^* = \arg\sup(\tilde{y}(u)^2) = u_0^* + \frac{\kappa'(u_0^*)}{|\kappa_0''(u_0^*)|}\frac{z_1}{z_0}. \tag{18}$$

Notice that it is dependent upon z_0, z_1 so that KS is no longer exactly the absolute value of a Gaussian. However it can be shown (after lenghty but straightforward calculation) that, to second order in ϵ, $\tilde{y}(u^*)$ *remains Gaussian*, albeit with a new width

$$\kappa^* \approx \sqrt{\kappa_0^2 + \kappa^2} > \kappa_0, \tag{19}$$

where all the functions are evaluated at u_0^*. In fact, this approximation works also with more than two modes, provided

$$\kappa(u)^2 \equiv \sum_{j\neq 0} \lambda_j U_j(u)^2 \ll \kappa_0(u)^2 = \lambda_0 U_0(u)^2, \tag{20}$$

in which case:

$$\kappa^* \approx \sqrt{\sum_j \lambda_j U_j(u_0^*)}. \tag{21}$$

3. An explicit example: The log-normal volatility model

In order to illustrate the above general formalism, we focus on the specific example of the product random variable $X = \sigma\xi$, with iid Gaussian residuals ξ and log-normal stochastic standard-deviations $\sigma = \mathrm{e}^{\omega}$. Such models are common in finance to describe stock returns, as will be discussed in the next section. For the time being, let us consider the case where the ω's are of zero mean, unit variance, and covariance given by:

$$\mathrm{Cov}(\omega_n\omega_{n+t}) = \Sigma^2 f\left(\frac{t}{T}\right), \qquad (t > 0). \tag{22}$$

The pairwise copulas in the covariance of $\overline{Y}$ can be explicitely written in the limit of weak correlations, $\Sigma^2 \to 0$. One finds:

$$C_t(u,v) - uv = \Sigma^2 f\left(\frac{t}{T}\right)\tilde{A}(u)\tilde{A}(v) \tag{23}$$

$$\text{with} \quad \tilde{A}(u) \;=\; \int_{-\infty}^{\infty} \varphi(\omega)\varphi'(\frac{F^{-1}(u)}{\mathrm{e}^{\omega}})\mathrm{d}\omega \tag{24}$$

where here and in the following $\varphi(\cdot)$ denotes the univariate Gaussian pdf, and $\Phi(\cdot)$ the Gaussian cdf. The spectrum of $A(u,v) = \tilde{A}(u)\tilde{A}(v)$ consists in a single non-degenerate eigenvalue $\lambda^A = \mathrm{Tr}A = \int_0^1 \tilde{A}(u)^2\,du = 1.176\cdot 10^{-2}$, and an infinitely degenerate null eigenspace. Assuming short-ranged memory, such that $f_\infty = \sum_{r=1}^{\infty} f(r) < +\infty$, the covariance kernel reads:

$$H(u,v) = I(u,v) + 2T\Sigma^2 f_\infty\, A(u,v).$$

Depending on the value of the parameters, the first term or the second term may be dominant. Note that one can be in the case of weak correlations ($\Sigma^2 \to 0$) but long range

memory $T \gg 1$, such that the product $T\Sigma^2$ can be large (this is the case of financial time series, see below). If $T\Sigma^2$ is small, one can use perturbation theory around the Brownian bridge limit (note that $\mathrm{Tr}I \approx 10\mathrm{Tr}A$, see Appendix), whereas if $T\Sigma^2$ is large, it is rather the Brownian term $I(u,v)$ that can be treated as a perturbation. Elements of the algebra necessary to set up these perturbation theories are given in the Appendix.

It is interesting to generalize the above model to account for weak dependence between the residuals ξ and between the residual and the volatility, without spoiling the log-normal structure of the model. We therefore write:

$$X_0 = \xi_0 \mathrm{e}^{\omega_0}; \quad X_t = \xi_t \mathrm{e}^{\alpha_t\omega_0 + \beta_t\xi_0 + \sqrt{1-\alpha_t^2-\beta_t^2}\omega_t} \quad \text{with} \quad \mathbb{E}[\xi_0\xi_t] = r_t$$

where all the variables are $\mathcal{N}(0,1)$, so that in particular

$$\begin{aligned}
\rho_t = \mathrm{Corr}(X_0, X_t) &= r_t(1+\beta_t^2)\mathrm{e}^{\alpha_t - 1} \\
\mathrm{Corr}(X_0^2, X_t^2) &= \frac{(1 + 2r_t^2(1 + 10\beta_t^2 + 8\beta_t^4) + 4\beta_t^2)\,\mathrm{e}^{4\alpha_t} - 1}{3\mathrm{e}^4 - 1} \\
\mathrm{Corr}(X_0, X_t^2) &= 2\beta_t \frac{(1 + 2r_t^2(1 + 2\beta_t^2))\,\mathrm{e}^{2\alpha_t - \frac{1}{2}}}{\sqrt{\mathrm{e}^4 - 1}}
\end{aligned}$$

The univariate marginal distributions of X_0 and X_t are identical and their cdf is given by the integral

$$F(x) = \int_{-\infty}^{\infty} \varphi(\omega)\Phi(\frac{x}{\mathrm{e}^\omega})\mathrm{d}\omega. \tag{25}$$

Expanding the bivariate cdf (or the copula) in the small dependence parameters $\alpha_t, \beta_t, \rho_t$ around $(0,0,0)$, we get

$$\begin{aligned}
C_t(u,v) - uv &\approx \alpha_t A(u,v) - \beta_t B(u,v) + \rho_t R(u,v) \\
&\approx \alpha_t \tilde{A}(u)\tilde{A}(v) - \beta_t \tilde{R}(u)\tilde{A}(v) + \rho_t \tilde{R}(u)\tilde{R}(v)
\end{aligned} \tag{26}$$

where $\tilde{A}(u)$ was defined above in Equation (24), and

$$\tilde{R}(u) = \int_{-\infty}^{\infty} \varphi(\omega)\varphi(\frac{F^{-1}(u)}{\mathrm{e}^\omega})\mathrm{d}\omega = \tilde{R}(1-u). \tag{27}$$

The contributions of $A(u,v), B(u,v)$ and $R(u,v)$ on the diagonal are illustrated in Figure 1. Notice that the term $B(u,v)$ (coming from cross-correlations between ξ_0 and ω_t, i.e. the so-called leverage effect, see below) breaks the symmetry $C_t(u,v) \neq C_t(v,u)$.

We now turn to a numerical illustration of our theory, in the simple case where only volatility correlations are present (i.e. $\beta_t = \rho_t = 0$ in Equation (26) above). We furthermore assume a Markovian dynamics for the log-volatilities:

$$X_n = \xi_n \mathrm{e}^{\omega_n - \mathbb{V}[\omega]}, \quad \text{with} \quad \omega_{n+1} = g\omega_n + \Sigma\eta_n, \tag{28}$$

where $g < 1$ and η_n are iid Gaussian variables of zero mean and unit variance. In this case,

$$\alpha_t = \ \mathrm{Cov}(\omega_n\omega_{n+t}) = \frac{\Sigma^2}{1-g^2}\,g^t. \tag{29}$$

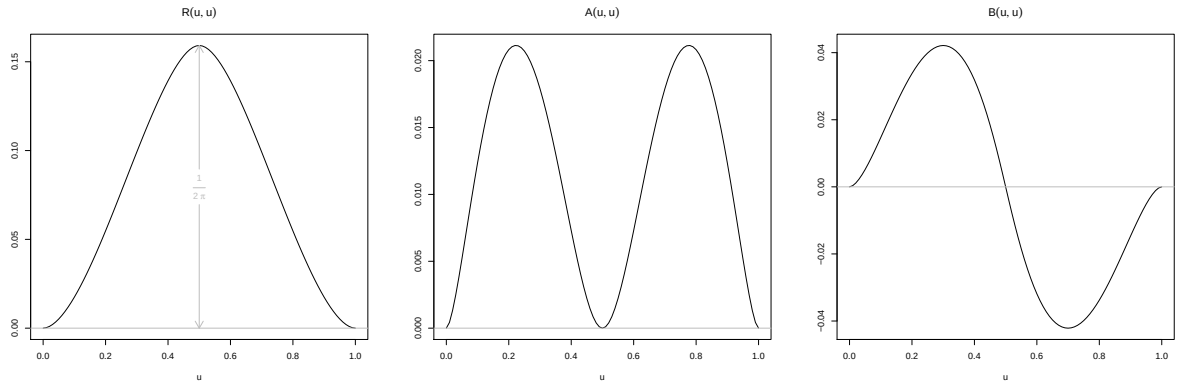

Figure 1. Copula diagonal of the log-normal volatility model: linear corrections to independence. **Left:** correction $R(u,u)$ due to correlation of the residuals (vertical axis in multiples of ρ) **Middle:** correction $A(u,u)$ due to correlation of the log-vols (vertical axis in multiples of α) **Right:** correction $B(u,u)$ due to leverage effect (vertical axis in multiples of $-\beta$)

In the limit where $\Sigma^2 \ll 1$, the weak dependence expansion holds and one finds explicitly:

$$H(u,v) = I(u,v) + 2\frac{g\Sigma^2}{(1-g)^2(1+g)}\,A(u,v). \tag{30}$$

In order to find the limit distribution of the test statistics, we procede by Monte-Carlo simulations. The range $[0,1]^2$ of the copula is discretized on a regular lattice of size $(M \times M)$. The limit process is described as a vector with M components and built from Equation (9) as $\tilde{\mathbf{y}} = U\Lambda^{\frac{1}{2}}\mathbf{z}$ where the diagonal elements of Λ are the eigenvalues of H (in decreasing order), and the columns of U are the corresponding eigenvectors. Clearly, $\mathrm{Cov}(\tilde{\mathbf{y}},\tilde{\mathbf{y}}) = U\Lambda U^\dagger = H$.

For each Monte-Carlo trial, M independent random values are drawn from a standard Gaussian distribution and collected in $\mathbf{z}$. Then $\mathbf{y}$ is computed using the above representation. This allows one to determine the two relevant statistics:

$$\begin{aligned} KS &= \max_{u=1\ldots M} |\tilde{y}_u| \\ CM &= \frac{1}{M}\sum_{u=1}^{M} \tilde{y}_u^2 = \frac{1}{M}\tilde{\mathbf{y}}^\dagger\tilde{\mathbf{y}} = \frac{1}{M}\mathbf{z}^\dagger\Lambda\mathbf{z}. \end{aligned}$$

The empirical cumulative distribution functions of the statistics for a large number of trials are shown in Figure 2 together with the usual Kolmogorov-Smirnov and Cramér-von-Mises limit distributions corresponding to the case of independent variables.

In order to check the accuracy of the obtained limit distribution, we generate 350 series of $N = 2500$ dates according to Equation (28). For each such series, we perform two GoF tests, namely KS and CM, and calculate the corresponding p-values. By construction, the p-values of a test should be uniformly distributed if the underlying distribution of the simulated data is indeed the same as the hypothesized distribution. In our case, when using the usual KS and CM distributions for independent data, the p-values are much too small and their histogram is statistically not compatible with the

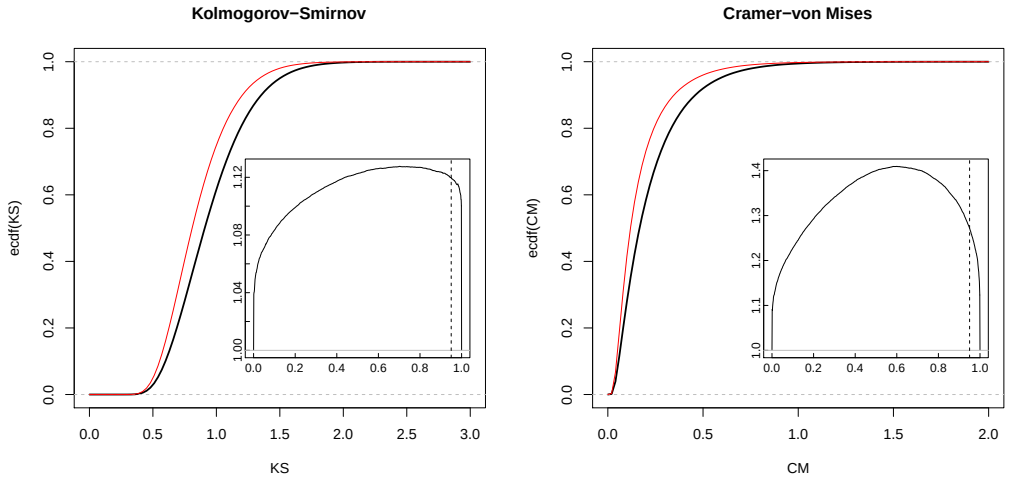

Figure 2. Markovian model – **Left:** Cumulative distribution function of the supremum of $\tilde{y}(u)$. **Right:** Cumulative distribution of the norm-2 of $\tilde{y}(u)$. The cases of independant drawings (thin red) and dependant drawings (bold black) are compared. The dependent observations are drawn according to the weak-dependence kernel (30) with parameters $g = 0.88, \Sigma^2 = 0.05$. **Insets:** The effective reduction ratio $\sqrt{\frac{N}{N_{\text{eff}}(u)}} = \frac{\text{ecdf}^{-1}(u)}{\text{cdf}_{\text{L}}^{-1}(u)}$ where L = KS,CM. The dashed vertical line is located at the 95-th centile and thus indicates the reduction ratio corresponding to the p-value $p = 0.05$ (as the test is unilateral).

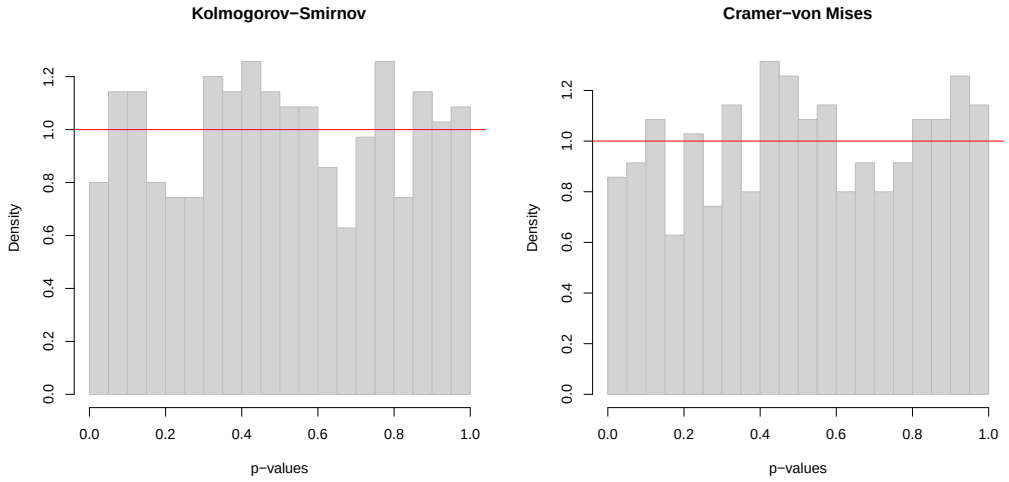

Figure 3. Histogram of the p-values in the GoF test on simulated data, according to Equation (28). Uniform distribution of the p-values of a test indicates that the correct law of the statistics is used.

uniform distribution. Instead, when using the appropriate limit distribution found by Monte-Carlo and corresponding to the correlation kernel (30), the calculated p-values are uniformly distributed, as can be visually seen on Figure 3, and as revealed statistically by a KS test (on the KS test!), comparing the 350 p-values to $\text{H}_0 : p \sim \mathcal{U}[0, 1]$.

If, instead of the AR(1) (Markovian) prescription (28), the dynamics of the ω_n is given by a Fractional Gaussian Noise (i.e. the increments of a fractional Brownian motion) [17] with Hurst index $\frac{2-\nu}{2} > \frac{1}{2}$, the log-volatility has a long ranged autocovariance

$$\alpha_t = \mathrm{Cov}(\omega_n, \omega_{n+t}) = \frac{\Sigma^2}{2}\left((t+1)^{2-\nu} - 2t^{2-\nu} + |t-1|^{2-\nu}\right), \quad t \geq 0 \tag{31}$$

that decays as a power law $\propto (2-3\nu+\nu^2)t^{-\nu}$ as $t \to \infty$, corresponding to long-memory, for which the above theory is not expected to be correct. The corresponding covariance kernel of the Xs,

$$H(u,v) = I(u,v) + 2\Sigma^2 A(u,v) \sum_{t=1}^{N} \left(1 - \frac{t}{N}\right) \alpha_t, \tag{32}$$

is used in a Monte-Carlo simulation like in the previous case in order to find the appropriate distribution of the test statistics KS and CM (shown in Figure 4, see caption for the choice of parameters). We again apply the GoF tests to simulated series, and compute the p-values according to the theory above. Because of the long-range dependence of the process, the p-values are not expected to be uniformly distributed any longer. However, the distribution of the p-values is still significantly corrected toward the uniform distribution, see Figure 4. It would be very interesting to extend the formalism above to this case as well.

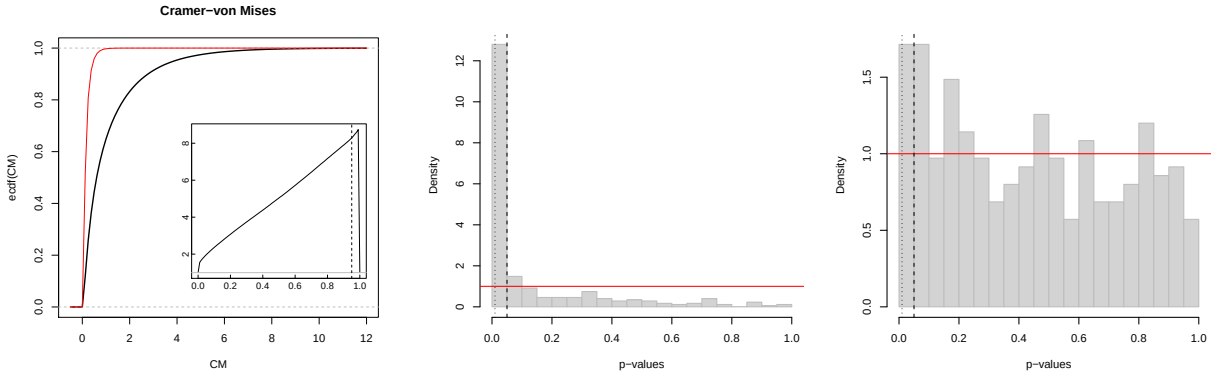

Figure 4. Fractional Brownian Motion – **Left:** Cumulative distribution function of the norm-2 of $\tilde{y}(u)$, see Fig. 2 for full caption. **Middle-Right:** Histogram of the p-values in the CM test on simulated data; when the naive CM distribution is used (middle), the obtained p-values are localized indicating that the wrong law is used. When we use our prediction for short-range dependencies (right), we find a clear improvement, as the p-values are more widely spread on $[0,1]$. However the p-values are still not statistically consistent with the uniform distribution. The dependent observations are drawn according to (31) with parameters $\nu = \frac{2}{5}, \Sigma^2 = 1, N = 1500$.

4. Application to financial time series

4.1. Stylized facts of daily stock returns

One of the contexts where long-ranged persistence is present is time series of financial asset returns. At the same time, the empirical determination of the distribution of these returns is of utmost importance, in particular for risk control and derivative pricing. As we will see, the volatility correlations are so long-ranged that the number of effectively independent observations is strongly reduced, in such a way that the GoF tests are not very tight, even with time series containing thousands of raw observations.

It is well-known that stock returns exhibit dependences of different kinds:

- at relatively high frequencies (up to a few days), returns show weak, but significant negative linear auto-correlations (see e.g. [18]);
- the absolute returns show persistence over very long periods, an effect called multiscale volatility clustering and for which several interesting models have been proposed in the last ten years [19, 20, 21, 22, 23, 24];
- past negative returns favor increased future volatility , an effect that goes under the name of "leverage correlations" in the literature [4, 25, 26, 27, 28, 29].

Our aim here is neither to investigate the origin of these effects and their possible explanations in terms of behavioral economics, nor to propose a new family of models to describe them. We rather want to propose a new way to characterize and measure the full structure of the temporal dependences of returns based on copulas, and extract from this knowledge the quantities needed to apply GoF tests to financial times series.

Throughout this section, the empirical results are based on a data set consisting of the daily returns of the stock price of listed large cap US companies. More precisely we keep only the 376 names present in the S&P-500 index constantly over the five years period 2000–2004, corresponding to $N = 1256$ days. The individual series are standardized, but this does not change the determination of copulas, that are invariant under increasing and continuous transformations of the marginals.

4.2. Empirical self-copulas

For each (u, v) on a lattice, we determine the lag dependent "self-copula" $C_t(u, v)$ by assuming stationarity, i.e. that the pairwise copula $C_{nm}(u, v)$ only depends on the time lag $t = m - n$. We also assume that all stocks are characterized by the same self-copula, and therefore an average over all the stock names in the universe is done in order to remove noise. This assumption could be refined, and some systematic effects of market cap, liquidity, tick size, etc, could be sought for, but we leave this for subsequent investigations.

The self-copulas are estimated non-parametrically with a bias correction, then fitted to the parametric family of log-normal copulas introduced in the previous section. We assume (and check a posteriori) that the weak dependence expansion holds, leaving us

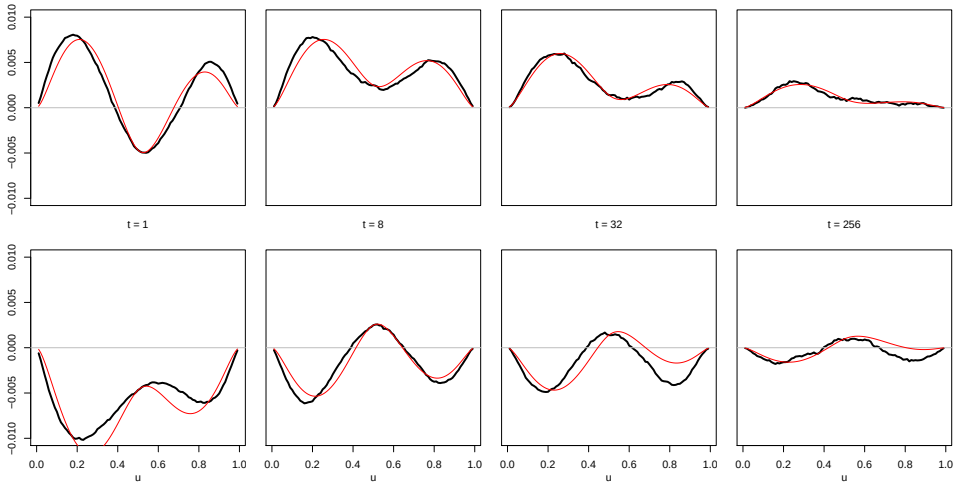

Figure 5. Diagonal (top) and anti-diagonal (bottom) of the self-copula for different lags; the product copula has been subtracted. A fit with Equation (26) is shown in thin red. Note that the y scale is small, confirming that the weak dependence expansion is justified. The dependence is still significant even for $t \sim 500$ days.

with three functions of time, α_t, β_t and ρ_t, to be determined. We fit for each t the copula diagonal $C_t(u,u)$ to Equation (26) above, determine α_t, β_t and ρ_t, and test for consistency on the anti-diagonal $C_t(u,1-u)$. Alternatively, we could determine these coefficients to best fit $C_t(u,v)$ in the whole (u,v) plane, but the final results are not very different. The results are shown in Figure 5 for lags $t = 1, 8, 32, 256$ days. Fits of similar quality are obtained up to $t = 512$.

Before discussing the time dependence of the fitted coefficients α_t, β_t and ρ_t, let us describe how the different effects show up in the plots of the diagonal and anti-diagonal copulas. The contribution of the linear auto-correlation can be directly observed at the central point $C_t(\frac{1}{2},\frac{1}{2})$ of the copula. It is indeed known [9] that for any pseudo-elliptical model (including the present log-normal framework) one has:

$$C_t(\tfrac{1}{2},\tfrac{1}{2}) = \frac{1}{4} + \frac{1}{2\pi}\arcsin\rho_t.$$

Note that this relation holds beyond the weak dependence regime. If $\beta_t^{(\mathrm{B})} = C_t(\frac{1}{2},\frac{1}{2}) - \frac{1}{4}$ — this is in fact Blomqvist's beta coefficient [30] — the auto-correlation is measured by $\rho_t = \sin(2\pi\beta_t^{(\mathrm{B})})$.

The volatility clustering effect can be visualized in terms of the diagonals of the self-copula; indeed, the excess (unconditional) probability of large events following previous large events of the same sign is $(C(u,u) - u^2)$ with $u < \frac{1}{2}$ for negative returns, and $u > \frac{1}{2}$ for positive ones. On the anti-diagonal, the excess (unconditional) probability of large positive events following large negative ones is, for small $u < \frac{1}{2}$, the upper-left volume $(C(u,1) - u\cdot 1) - (C(u,1-u) - u(1-u)) = u(1-u) - C(u,1-u)$ and similarly the excess probability of large negative events following large positive ones is the same expression for large $u > \frac{1}{2}$ (lower-right volume). As illustrated on Figure 5, these four quadrants exceed the independent case prediction, suggesting a genuine clustering of

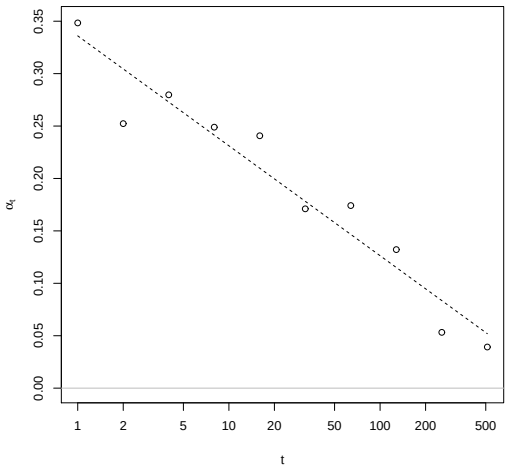

Figure 6. Auto-correlation of the volatilities, for lags ranging from 1 to 768 days. Each point represents the value of α_t extracted from a fit of the empirical copula diagonal at a given lag to the relation (26). We also show the fit to a multifractal model, $\alpha_t = -\Sigma^2 \log \frac{t}{T}$, with $\Sigma^2 = 0.046$ and $T = 1467$ days.

the amplitudes, measured by α_t. Finally, an asymmetry is clearly present: the effect of large negative events on future amplitudes is stronger than the effect of previous positive events. This is an evidence for the leverage effect: negative returns cause a large volatility, which in turn makes future events (positive or negative) to be likely larger. This effect is captured by the coefficient β_t.

The evolution of the coefficients α_t, β_t and ρ_t for different lags reveals the following properties: i) the linear auto-correlation ρ_t is short-ranged (a few days), and negative; ii) the leverage parameter β_t is short-ranged and, as is well known, negative, revealing the asymmetry discussed above; iii) the correlation of volatility is *long-ranged* and of relatively large positive amplitude (see Figure 6), in line with the known long range volatility clustering. More quantitatively, we find that the parameter α_t for lags ranging from 1 to 768 days is consistent with an effective relation well fitted by the "multifractal" [31, 20, 32, 21] prediction for the volatility autocorrelations: $\alpha_t = -\Sigma^2 \log \frac{t}{T}$, with an amplitude $\Sigma^2 = 0.046$ and a horizon $T = 1467$ days consistent, in order of magnitude, with previous determinations.

The remarkable point, already noticed in previous empirical works on multifractals [31, 33], is that the horizon T, beyond which the volatility correlations vanish, is found to be extremely long. In fact, the extrapolated horizon T is larger than the number of points of our sample N! This long correlation time has two consequences: first, the parameter $2T\Sigma^2 f_\infty$ that appears in the kernel $H(u, v)$ is large, ≈ 135. This means that the dependence part $T\Sigma^2 f_\infty A(u, v)$ is dominant over the independent Brownian bridge part $I(u, v)$. This is illustrated in Figure 7, where we show the first eigenvector of $H(u, v)$, which we compare to the non-zero eigenmode of $A(u, v)$, and to the first eigenvector of $I(u, v)$. Second, the hypothesis of a stationnary process, which requires

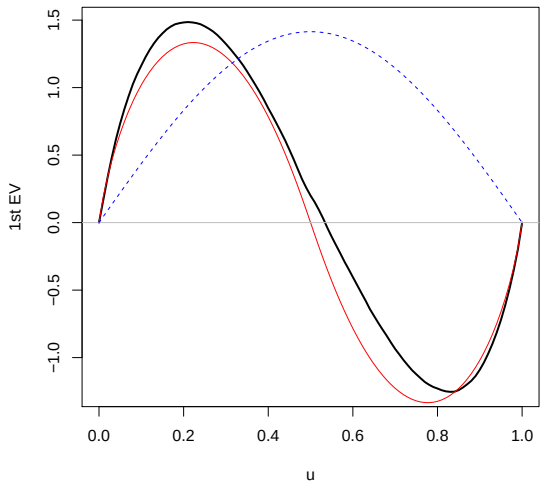

Figure 7. Bold black: The first eigenvector of the empirical kernel $H(u,v) = I(u,v)\,[1+\Psi_N(u,v)]$. **Plain red:** The function $\tilde{A}(u)$ (normalized), corresponding to the pure effect of volatility clustering in a log-normal model, in the limit where the Brownian bridge contribution $I(u,v)$ becomes negligible. **Dashed blue:** The largest eigenmode $|1\rangle = \sqrt{2}\sin(\pi u)$ of the independent kernel $I(u,v)$.

that $N \ll T$, is not met here, so we expect important preasymptotic corrections to the above theoretical results.

4.3. Monte-Carlo estimation of the limit distributions

Since $H(u,v)$ is copula-dependent, and considering the poor analytical progress made about the limit distributions of *KS* and *CM* in cases other than independence, the asymptotic laws will be computed numerically by Monte-Carlo simulations (like in the example of Section 3) with the empirically determined $H(u,v)$.

The empirical cumulative distribution functions of the statistics for a large number of trials are shown in Figure 8 together with the usual Kolmogorov-Smirnov and Cramér-von-Mises limit distributions corresponding to the case of independent variables. One sees that the statistics adapted to account for dependences are stretched to the right, meaning that they accept higher values of *KS* or *CM* (i.e. measures of the difference between the true and the empirical distributions). In other words, the outcome of a test based on the usual *KS* or *CM* distributions is much more likely to be negative, as it will consider "high" values (like 2–3) as extremely improbable, whereas a test that accounts for the strong dependence in the time series would still acccept the null-hypothesis for such values.

As an illustration, we apply the test of Cramér-von Mises to our dataset, comparing the empirical univariate distributions of stock returns to a simple model of log-normal

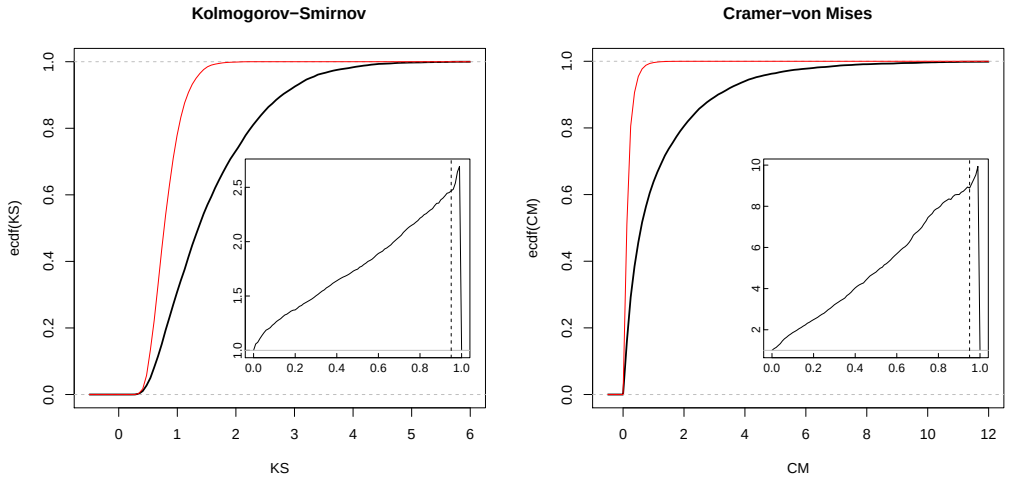

Figure 8. Left: Cumulative distribution function of the supremum of $\tilde{y}(u)$. **Right:** Cumulative distribution of the norm-2 of $\tilde{y}(u)$. The cases of independant drawings (thin red) and dependant drawings (bold black) are compared. The dependent observations are drawn according to the empirical average self-copula of US stock returns in 2000-2004. **Insets:** The effective reduction ratio $\sqrt{\frac{N}{N_{\text{eff}}(u)}} = \frac{\text{ecdf}^{-1}(u)}{\text{cdf}_{\text{L}}^{-1}(u)}$ where L = KS,CM.

stochastic volatility

$$X = \mathrm{e}^{s\omega - s^2}\xi \qquad \text{where} \quad \xi, \omega \overset{\text{iid}}{\sim} \mathcal{N}(0,1). \tag{33}$$

The volatility of volatility parameter s can be calibrated from the time series $\{x_t\}_t$ as

$$s^2 = \log\left(\frac{2}{\pi}\frac{\langle x_t^2\rangle_t}{\langle x_t\rangle_t^2}\right). \tag{34}$$

We want to test the hypothesis that the log-normal model with a unique value of s *for all stocks* is compatible with the data. In practice, for each stock i, s_i is chosen as the average of (34) over all *other* stocks in order to avoid endogeneity issues and biases in the calculations of the p-values of the test. s_i is found to be ≈ 0.5 and indeed almost identical for all stocks. Then the GoF statistic *CM* is computed for each stock i and the corresponding p-value is calculated.

Figure 9 shows the distribution of the p-values, as obtained by using the usual asymptotic Cramér-von Mises distribution for independent samples (left) and the modified version allowing for dependence (right). We clearly observe that the standard Cramér-von Mises test strongly rejects the hypothesis of a common log-normal model, as the corresponding p-values are strongly concentrated around zero, which leads to an excessively high rejection rate. The use of the generalized Cramér-von Mises test for dependent variables greatly improves the situation, with in fact now too many high values of p. The hypothesis of a common log-normal model for all stocks cannot be rejected when the long-memory of volatility is taken into account. The overabundant large values of p may be due to the fact that all stocks are in fact exposed to a common volatility factor (the "market mode"), which makes the estimation of s somewhat

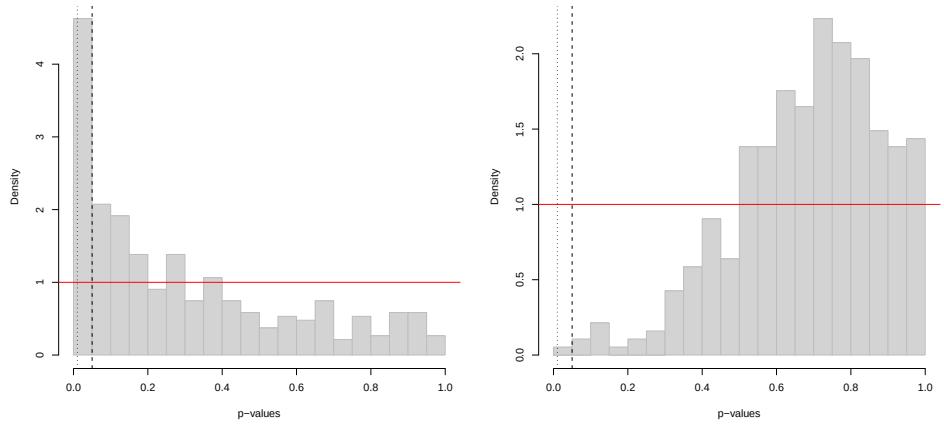

Figure 9. Histogram of the p-values in a Cramér-von Mises-like test, see the text for the test design. **Left:** when using the law of Cramér-von Mises, the obtained p-values are far from uniformly distributed and strongly localized under the threshold $p = 0.05$ (dashed vertical line) occasioning numerous spurious rejections. **Right:** when using the modified law taking dependences into account, the test rejects the hypothesis of an identical distribution for all stocks much less often.

endogeneous and generates an important bias. Another reason is that the hypothesis that the size of the sample N is much larger than the correlation time T does not hold for our sample, and corrections to our theoretical results are expected in that case.§ It would be actually quite interesting to extend the above formalism to the long-memory case, where $T \gg N \gg 1$.

§ Note that in practice, we have estimated $\Psi_N(u, v)$ by summing the empirically determined copulas up to $t_{\max} = 512$, which clearly underestimates the contribution of large lags.

5. Conclusion

The objectives of this paper were twofold: on the theoretical side, we introduced a framework for the study of statistical tests of Goodness-of-Fit with dependent samples; on the empirical side, we presented new measurement as well as phenomenological models for non-linear dependences in time series of daily stock returns. Both parts heavily rely on the notion of bivariate self-copulas.

In summary, GoF testing on persistent series cannot be universal as is the case for iid variables, but requires a careful estimation of the self-copula at all lags. Correct asymptotic laws for the Kolmogorov-Smirnov and Cramér-von Mises statistics can be found as long as dependences are short ranged, i.e. $T \ll N$. From the empirical estimation of the self-copula of US stock returns, long-ranged volatility clustering with multifractal properties is observed as the dominant contribution to self-dependence, in line with previous studies. However, subdominant modes are present as well and a precise understanding of those involves an in-depth study of the spectral properties of the correlation kernel H.

One of the remarkable consequence of the long-memory nature of the volatility is that the number of effectively independent observations is significantly reduced, as both the Kolmogorov-Smirnov and Cramér-von Mises tests accept much larger values of the deviations (see Figure 8). As a consequence, it is much more difficult to reject the adequation between a reasonable statistical model and empirical data. We suspect that many GoF tests used in the literature to test models of financial returns are fundamentally flawed because of the long-ranged volatility correlations. In intuitive terms, the long-memory nature of the volatility can be thought of as a sequence of volatility regime shifts, each with a different lifetime, and with a broad distribution of these lifetimes. It is clear that in order to fully sample the unconditional distribution of returns, all the regimes must be encountered several times. In the presence of volatility persistence, therefore, the GoF tests are much less stringent, because there is always a possibility that one of these regimes was not, or only partially, sampled.

We conclude with two remarks of methodological interest.

1) The method presented for dealing with self-dependences while using statistical tests of Goodness-of-Fit is computationally intensive in the sense that it requires to estimate empirically the self-copula for all lags over the entire unit square. In the non-parametric setup, discretization of the space must be chosen so as to provide a good approximation of the continuous distance measures while at the same time not cause too heavy computations. Considering that fact, it is often more appropriate to use the Cramér-von Mises-like test rather than the Kolmogorov-Smirnov-like, as numerical error on the evaluation of the integral will typically be much smaller than on the evaluation of the supremum on a grid, more so when the grid size is only about $\frac{1}{M} \approx \frac{1}{100}$.

2) The case with long-ranged dependence $T \gg N \gg 1$ cannot be treated in the framework presented here. First because the Central Limit Theorem does not hold in that case, and finding the limit law of the statistics may require more advanced

mathematics. But even pre-asymptotically, summing the lags over the available data up to $t \approx N$ means that a lot of noise is included in the determination of Ψ_N. This, in turn, is likely to cause the empirically determined kernel H not to be positive definite. One way of addressing this issue is to follow a semi-parametric procedure: the copula C_t is still estimated non-parametrically, but the kernel H sums the lagged copulas C_t only up to a scale where the linear correlations and leverage correlations vanish, and only one long-ranged dependence mode remains. This last contribution can be fitted by an analytical form, that can then be summed up to its own scale, or even to infinity.

In terms of financial developments, we believe that an empirical exploration of the self-copulas for series of diverse asset returns and at diverse frequencies is of primordial importance in order to grasp the complexity of the non-linear time-dependences. In particular, expanding the concept of the self-copula to pairs of assets is likely to reveal subtle dependence patterns. From a practitioner's point of view, a multivariate generalization of the self-copula could lead to important progresses on such issues as causality, lead-lag effects and the accuracy of multivariate prediction.

Acknowledgments

We thank Frédéric Abergel for helpful comments and Vincent Vargas for fruitful discussions.

Appendix A. Pseudo-elliptical copula: expansion around independence

We compute here the spectrum and eigenvectors of the kernel $H(u,v)$ in the case of pseudo-elliptical copula with weak dependences, starting from the expansion (26).

The situation is better understood in terms of operators acting in the Hilbert space of continuous functions on $[0,1]$ vanishing in the border. Using Dirac's braket notations, $A = |\tilde{A}\rangle\langle\tilde{A}|$, $B = |\tilde{R}\rangle\langle\tilde{A}|$, $R = |\tilde{R}\rangle\langle\tilde{R}|$. The sine functions $|j\rangle = \sqrt{2}\sin(j\pi u)$ build a basis of this Hilbert space, and interestingly they are the eigenvectors of the independent kernel $I(u,v)$ (I stands for 'Independence' and is the covariance matrix of the Brownian motion: $I = M - P$ where M denotes the bivariate upper Fréchet-Hoeffding copula and P the bivariate product copula).

It is then easy to find the spectra: rank-one operators have at most one non-null eigenvalue. Using the parities of $\tilde{A}(u)$ and $\tilde{R}(u)$ with respect to $\frac{1}{2}$ and imposing orthonormality of the eigenvectors, we can sketch the following table of the non zero eigenvalues and eigenvectors of the different operators:

$$\begin{array}{ll} \lambda^I_j = (j\pi)^{-2} & U^I_j(u) \;=\; |j\rangle \\ \lambda^R = \langle\tilde{R}|\tilde{R}\rangle = \mathrm{Tr}R & |U^R_0\rangle \;=\; |\tilde{R}\rangle/\sqrt{\mathrm{Tr}R} \\ \lambda^A = \langle\tilde{A}|\tilde{A}\rangle = \mathrm{Tr}A & |U^A_0\rangle \;=\; |\tilde{A}\rangle/\sqrt{\mathrm{Tr}A} \end{array}$$

For the pseudo-elliptical copula with weak dependence, H has the following general form:

$$H = I + \tilde{\rho}R + \tilde{\alpha}A - \frac{\tilde{\beta}}{2}(B + B^\dagger). \tag{A.1}$$

The operator $B + B^\dagger$ has two non zero eigenvalues $\pm\sqrt{\lambda^R\lambda^A}$, with eigenvectors $(|U^R_0\rangle \pm |U^A_0\rangle)/\sqrt{2}$. In order to approximately diagonalize H, it is useful to notice that in the present context A and R are close to commuting with I. More precisely, it turns out that $|U^A_0\rangle$ is very close to $|2\rangle$, and $|U^R_0\rangle$ even closer to $|1\rangle$. Indeed, $a_2 = \langle U^A_0|2\rangle \approx 0.9934$ and $r_1 = \langle U^R_0|1\rangle \approx 0.9998$. Using the symmetry of A and R, we can therefore write:

$$\begin{array}{lll} |U^A_0\rangle = a_2|2\rangle + \epsilon_a|2_\perp\rangle & \text{with} & \langle 2|2_\perp\rangle = \langle 2j-1|2_\perp\rangle = 0, \forall j \geq 1 \\ |U^R_0\rangle = r_1|1\rangle + \epsilon_r|1_\perp\rangle & \text{with} & \langle 1|1_\perp\rangle = \langle 2j|1_\perp\rangle = 0, \forall j \geq 1 \end{array}$$

where $\epsilon_a = \sqrt{1-a_2^2} \ll 1$ and $\epsilon_r = \sqrt{1-r_1^2} \ll 1$. The components of $|2_\perp\rangle$ on the even eigenvectors of I are determined as:

$$\langle 2_\perp|2j\rangle = \frac{\langle U^A_0|2j\rangle}{\epsilon_a} \qquad j \geq 2,$$

Table A1. Traces of the operators appearing in the covariance functions (multiples of 10^{-2}). Traces of the powers of the rank-one A, R equal powers of their traces. The trace of $B + B^\dagger$ is zero.

I	A	R	I^2	IA	IR
16.667	1.176	7.806	111.139	2.948	79.067

and similarly:

$$\langle 1_\perp|2j-1\rangle = \frac{\langle U_0^R|2j-1\rangle}{\epsilon_r} \qquad j \geq 2.$$

Using the definition of the coefficients α_t, β_t and ρ_t given in section 3, we introduce the following notations:

$$\tilde{\alpha} = 2\,\mathrm{Tr}A \lim_{N\to\infty} \sum_{t=1}^{N-1} \left(1 - \frac{t}{N}\right) \alpha_t$$

$$\tilde{\rho} = 2\,\mathrm{Tr}R \lim_{N\to\infty} \sum_{t=1}^{N-1} \left(1 - \frac{t}{N}\right) \rho_t$$

$$\tilde{\beta} = 2\sqrt{\mathrm{Tr}A\,\mathrm{Tr}R} \lim_{N\to\infty} \sum_{t=1}^{N-1} \left(1 - \frac{t}{N}\right) \beta_t$$

so that H writes:

$$\begin{aligned} H = I \quad & + \tilde{\alpha}|U_0^A\rangle\langle U_0^A| + \tilde{\rho}|U_0^R\rangle\langle U_0^R| - \tilde{\beta}\,|U_0^R\rangle\!\!\overleftrightarrow{\rangle}\langle U_0^A| \\ = H_0 & + \epsilon_a \left(\tilde{\alpha} a_2\,|2\rangle\!\!\overleftrightarrow{\langle} 2_\perp| - \tilde{\beta} r_1 a_\perp\,|1\rangle\!\!\overleftrightarrow{\langle} 2_\perp|\right) \\ & + \epsilon_r \left(\tilde{\rho} r_1\,|1_\perp\rangle\!\!\overleftrightarrow{\langle} 1| - \tilde{\beta} a_2\,|1_\perp\rangle\!\!\overleftrightarrow{\langle} 2|\right) \\ & + \left(\tilde{\alpha}\epsilon_a^2|2_\perp\rangle\langle 2_\perp| + \tilde{\rho}\epsilon_r^2|1_\perp\rangle\langle 1_\perp| - \tilde{\beta}\epsilon_a\epsilon_r\,|2_\perp\rangle\!\!\overleftrightarrow{\rangle}\langle 1_\perp|\right) \end{aligned}$$

where $|\psi_1\rangle\!\overleftrightarrow{\rangle}\langle\psi_2| = \frac{1}{2}(|\psi_1\rangle\langle\psi_2| + |\psi_2\rangle\langle\psi_1|)$ and H_0 is the unperturbed operator (0-th order in both ϵs)

$$H_0 = \sum_{j\geq 3} \lambda_j^I |j\rangle\langle j| + (\lambda_2^I + \tilde{\alpha}a_2^2)|2\rangle\langle 2| + (\lambda_1^I + \tilde{\rho}r_1^2)|1\rangle\langle 1| - \tilde{\beta} r_1 a_2\,|2\rangle\!\!\overleftrightarrow{\rangle}\langle 1|$$

the spectrum of which is easy to determine as:

$$\begin{aligned} &\lambda_1^{H_0} = \lambda_- \overset{\tilde{\rho},\tilde{\beta}\to 0}{\longrightarrow} \lambda_1^I && |U_1^{H_0}\rangle = -\frac{|-\rangle}{\sqrt{\langle -|-\rangle}} \overset{\tilde{\rho},\tilde{\beta}\to 0}{\longrightarrow} |1\rangle \\ &\lambda_2^{H_0} = \lambda_+ \overset{\tilde{\rho},\tilde{\beta}\to 0}{\longrightarrow} \lambda_2^I + \tilde{\alpha}a_2^2 && |U_2^{H_0}\rangle = \frac{|+\rangle}{\sqrt{\langle +|+\rangle}} \overset{\tilde{\rho},\tilde{\beta}\to 0}{\longrightarrow} |2\rangle \\ &\lambda_j^{H_0} = \lambda_j^I && |U_j^{H_0}\rangle = |j\rangle \qquad (j \geq 3) \end{aligned}$$

where

$$\lambda_\pm = \frac{\lambda_1^I + \tilde{\rho}r_1^2 + \lambda_2^I + \tilde{\alpha}a_2^2 \pm \sqrt{(\lambda_1^I + \tilde{\rho}r_1^2 - \lambda_2^I - \tilde{\alpha}a_2^2)^2 + 4(\tilde{\beta}r_1 a_2)^2}}{2}$$

and $|\pm\rangle$ the corresponding eigenvectors, which are linear combination of $|1\rangle$ and $|2\rangle$ only. Therefore, $\langle 1_\perp|\pm\rangle = \langle 2_\perp|\pm\rangle = 0$. This implies that there is no corrections to the eigenvalues of H_0 to first order in the ϵs.

At the next order, instead, some corrections appear. We call:

$$\begin{aligned} V_{i,j} = & (\tilde{\rho} r_1 \langle 1|U_i^{H_0}\rangle - \frac{\tilde{\beta} a_2}{2}\langle 2|U_i^{H_0}\rangle)\langle j|1_\perp\rangle\epsilon_r \\ & + (\tilde{\alpha} a_2 \langle 2|U_i^{H_0}\rangle - \frac{\tilde{\beta} r_1}{2}\langle 1|U_i^{H_0}\rangle)\langle j|2_\perp\rangle\epsilon_a \end{aligned}$$

the matrix elements of the first order perturbation of H, whence

$$\lambda_1^H = \lambda_1^{H_0} + \sum_{j\geq 3} \frac{V_{1,j}^2}{\lambda_1^{H_0} - \lambda_j^{H_0}}$$
$$\lambda_2^H = \lambda_2^{H_0} + \sum_{j\geq 3} \frac{V_{2,j}^2}{\lambda_2^{H_0} - \lambda_j^{H_0}}$$
$$\lambda_j^H = \lambda_j^{H_0} + \sum_{i=1,2} \frac{V_{i,j}^2}{\lambda_j^{H_0} - \lambda_i^{H_0}} + (\tilde{\alpha}\epsilon_a^2\langle j|2_\perp\rangle^2 + \tilde{\rho}\epsilon_r^2\langle j|1_\perp\rangle^2 - \tilde{\beta}\epsilon_a\epsilon_r\langle j|1_\perp\rangle\langle j|2_\perp\rangle)$$

As of the eigenvectors, it is enough to go to first order in ϵs to get a non-trivial perturbative correction:

$$|U_1^H\rangle = |U_1^{H_0}\rangle + \sum_{j\geq 3} \frac{V_{1,j}}{\lambda_1^{H_0} - \lambda_j^{H_0}}|j\rangle$$
$$|U_2^H\rangle = |U_2^{H_0}\rangle + \sum_{j\geq 3} \frac{V_{2,j}}{\lambda_2^{H_0} - \lambda_j^{H_0}}|j\rangle$$
$$|U_j^H\rangle = |j\rangle + \sum_{i=1,2} \frac{V_{i,j}}{\lambda_j^{H_0} - \lambda_i^{H_0}}|U_i^{H_0}\rangle$$

The special case treated numerically in section 3 corresponds to $\tilde{\rho} = \tilde{\beta} = 0$, such that the above expressions simplify considerably, since in that case $V_{1,j} \equiv 0$ and $V_{2,2j-1} = 0$, while $V_{2,2j} = \tilde{\alpha}a_2\langle U_0^A|2j\rangle$. To first order in the ϵs, the spectrum is not perturbed and calls $\lambda_i^H = \lambda_i^{H_0} = \lambda_i^I + \tilde{\alpha}a_2^2\delta_{i2}$, so that the characteristic function of the modified CM distribution is, according to Equation (12),

$$\phi(t) = \prod_j \left(1 - 2\mathrm{i}t/(j\pi)^2\right)^{-\frac{1}{2}} \times \sqrt{\frac{1 - 2\mathrm{i}t\lambda_2^I}{1 - 2\mathrm{i}t\lambda_2^{H_0}}}.$$

Its pdf is thus the convolution of the Fourier transform of $\phi_I(t)$ (characteristic function associated to the usual CM distribution [16]) and the Fourier transform of the correction $\phi_c(t) = \sqrt{1 - 2\mathrm{i}t\lambda_2^I}/\sqrt{1 - 2\mathrm{i}t\lambda_2^{H_0}}$. Noting that $(1-2\mathrm{i}\sigma^2 t)^{-\frac{1}{2}}$ is the characteristic function of the chi-2 distribution, it can be shown that for $k > 0$, and with $\mu \equiv \lambda_2^{H_0}$ for the sake of readability:

$$\begin{aligned}
\frac{1}{\sqrt{2\pi}}\mathrm{FT}\left(\phi_c\right) &= \delta(k) - \int_{\lambda_2^I}^{\mu} \mathrm{d}\lambda \frac{\partial}{\partial k}\left(\chi^2(k;\mu) * \chi^2(k;\lambda)\right) \\
&= \delta(k) - \int_{\lambda_2^I}^{\mu} \mathrm{d}\lambda \frac{\mathrm{e}^{-\frac{\lambda+\mu}{4\lambda\mu}k}}{8(\lambda\mu)^{\frac{3}{2}}}\left((\mu-\lambda)I_1(\frac{\mu-\lambda}{4\lambda\mu}k) - (\mu+\lambda)I_0(\frac{\mu-\lambda}{4\lambda\mu}k)\right) \\
&\approx \delta(k) + \mathrm{e}^{-\frac{k}{2\lambda}}\frac{\tilde{\alpha}a_2^2}{4\lambda^2}I_0(\frac{\tilde{\alpha}a_2^2}{4\lambda^2}k)
\end{aligned}$$

where $\chi^2(k;\sigma^2) = (2\pi\sigma^2 k\,\mathrm{e}^{k/\sigma^2})^{-\frac{1}{2}}$ is the pdf of the chi-2 distribution, I_n are the modified Bessel functions of the first kind, and $*$ denotes the convolution operation. The approximation on the last line holds as long as $\tilde{\alpha} \ll \lambda_2^I = (2\pi)^{-2}$ and in this regime

we obtain finally

$$\begin{aligned}\mathcal{P}[CM = k] &= \sqrt{2\pi}\mathrm{FT}(\phi)(k) = (\mathrm{FT}(\phi_I) * \mathrm{FT}(\phi_c))(k) \\ &= \mathcal{P}_I(k) + 4\tilde{\alpha}a_2^2\pi^4 \int_0^k \mathcal{P}_I(z)\mathrm{e}^{-2\pi^2(k-z)} I_0(4\tilde{\alpha}a_2^2\pi^4(k-z))\mathrm{d}z \\ &= \mathcal{P}_I(k) + 4\tilde{\alpha}a_2^2\pi^4 \int_0^k \mathcal{P}_I(k-z)\mathrm{e}^{-2\pi^2 z} I_0(4\tilde{\alpha}a_2^2\pi^4 z)\mathrm{d}z\end{aligned}$$

[1] D. A. Darling, "The Kolmogorov-Smirnov, Cramer-von Mises Tests," *The Annals of Mathematical Statistics*, vol. 28, no. 4, pp. 823–838, 1957.

[2] 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*, vol. 60, no. 6, p. 6519, 1999.

[3] A. A. Dragulescu and V. M. Yakovenko, "Probability distribution of returns in the heston model with stochastic volatility," *Quantitative Finance*, vol. 2, no. 6, pp. 443–453, 2002.

[4] J.-P. Bouchaud and M. Potters, "More stylized facts of financial markets: leverage effect and downside correlations," *Physica A: Statistical Mechanics and its Applications*, vol. 299, no. 1-2, pp. 60–70, 2001.

[5] Y. Malevergne and D. Sornette, *Extreme financial risks: From dependence to risk management*. Springer Verlag, 2006.

[6] P. Embrechts, F. Lindskog, and A. McNeil, "Modelling dependence with copulas and applications to risk management," *Handbook of heavy tailed distributions in finance*, vol. 8, pp. 329–384, 2003.

[7] P. Embrechts, A. McNeil, and D. Straumann, "Correlation and dependence in risk management: properties and pitfalls," *Risk management: value at risk and beyond*, pp. 176–223, 2002.

[8] T. Mikosch, "Copulas: Tales and facts," *Extremes*, vol. 9, no. 1, pp. 3–20, 2006.

[9] R. Chicheportiche and J.-P. Bouchaud, "The joint distribution of stock returns is not elliptical," *Arxiv preprint q-fin.ST/1009.1100*, 2010.

[10] B. K. Beare, "Copulas and temporal dependence," *Econometrica*, vol. 78, no. 1, pp. 395–410, 2010.

[11] R. Ibragimov and G. Lentzas, "Copulas and long memory." Harvard Institute of Economic Research discussion paper, 2008.

[12] A. J. Patton, "Copula-based models for financial time series," *Handbook of financial time series*, pp. 767–785, 2009.

[13] W. F. Darsow, B. Nguyen, and E. T. Olsen, "Copulas and markov processes," *Illinois Journal of Mathematics*, vol. 36, no. 4, pp. 600–642, 1992.

[14] R. C. Bradley, *Introduction to strong mixing conditions*, vol. 1–3. Kendrick Press, Heber City, Utah, 2007.

[15] X. Chen, L. P. Hansen, and M. Carrasco, "Nonlinearity and temporal dependence," *Journal of Econometrics*, vol. 155, no. 2, pp. 155–169, 2010.

[16] T. W. Anderson and D. A. Darling, "Asymptotic Theory of Certain "Goodness of Fit" Criteria Based on Stochastic Processes," *The Annals of Mathematical Statistics*, vol. 23, no. 2, pp. 193–212, 1952.

[17] B. B. Mandelbrot and J. W. Van Ness, "Fractional brownian motions, fractional noises and applications," *SIAM review*, vol. 10, no. 4, pp. 422–437, 1968.

[18] D. Avramov, T. Chordia, and A. Goyal, "Liquidity and autocorrelations in individual stock returns," *The Journal of Finance*, vol. 61, no. 5, pp. 2365–2394, 2006.

[19] M. Pasquini and M. Serva, "Multiscaling and clustering of volatility," *Physica A: Statistical Mechanics and its Applications*, vol. 269, no. 1, pp. 140–147, 1999.

[20] L. E. Calvet and A. Fisher, *Multifractal Volatility: Theory, Forecasting, and Pricing*. Academic Press, 2008.

[21] 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*, vol. 17, no. 3, pp. 537–548, 2000.

[22] G. Zumbach and P. Lynch, "Heterogeneous volatility cascade in financial markets," *Physica A: Statistical Mechanics and its Applications*, vol. 298, no. 3-4, pp. 521–529, 2001.

[23] P. Lynch and G. Zumbach, "Market heterogeneities and the causal structure of volatility," *Quantitative Finance*, vol. 3, no. 4, pp. 320–331, 2003.

[24] L. Borland, J.-P. Bouchaud, J.-F. Muzy, and G. Zumbach, "The dynamics of financial markets–mandelbrot's multifractal cascades, and beyond," *Wilmott Magazine*, pp. 86–96, March 2005.

[25] J. Perelló, J. Masoliver, and J.-P. Bouchaud, "Multiple time scales in volatility and leverage correlations: a stochastic volatility model," *Applied Mathematical Finance*, vol. 11, no. 1, pp. 27–50, 2004.

[26] B. Pochart and J.-P. Bouchaud, "The skewed multifractal random walk with applications to option smiles," *Quantitative Finance*, vol. 2, no. 4, pp. 303–314, 2002.

[27] Z. Eisler and J. Kertesz, "Multifractal model of asset returns with leverage effect," *Physica A: Statistical Mechanics and its Applications*, vol. 343, pp. 603–622, 2004.

[28] P. T. Ahlgren, M. H. Jensen, I. Simonsen, R. Donangelo, and K. Sneppen, "Frustration driven stock market dynamics: Leverage effect and asymmetry," *Physica A: Statistical Mechanics and its Applications*, vol. 383, no. 1, pp. 1–4, 2007.

[29] P.-A. Reigneron, R. Allez, and J.-P. Bouchaud, "Principal regression analysis and the index leverage effect," *Physica A: Statistical Mechanics and its Applications*, 2011.

[30] N. Blomqvist, "On a measure of dependence between two random variables," *The Annals of Mathematical Statistics*, vol. 21, no. 4, pp. 593–600, 1950.

[31] J.-F. Muzy, D. Sornette, J. Delour, and A. Arneodo, "Multifractal returns and hierarchical portfolio theory," *Quantitative Finance*, vol. 1, no. 1, pp. 131–148, 2001.

[32] 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*, vol. 26, p. 194, 2008.

[33] J. Duchon, R. Robert, and V. Vargas, "Forecasting volatility with the multifractal random walk model," *Mathematical Finance*, 2008.

[34] M. S. Weiss, "Modification of the Kolmogorov-Smirnov statistic for use with correlated data," *Journal of the American Statistical Association*, vol. 73, no. 364, pp. 872–875, 1978.