Correlation structure of extreme stock returns

Pierre Cizeau[†], Marc Potters[†] and Jean-Philippe Bouchaud[*,†]

[†]*Science & Finance*
The Research Division of Capital Fund Management
109–111 rue Victor-Hugo, 92532 Levallois CEDEX, *France*
http://www.science-finance.fr

[*] *Service de Physique de l'État Condensé,*
Centre d'études de Saclay,
Orme des Merisiers, 91191 Gif-sur-Yvette CEDEX, *FRANCE*

First version: June 2, 2000
This version
May 9, 2001

Abstract

It is commonly believed that the correlations between stock returns increase in high volatility periods. We investigate how much of these correlations can be explained within a simple non-Gaussian one-factor description with *time independent* correlations. Using surrogate data with the true market return as the dominant factor, we show that most of these correlations, measured by a variety of different indicators, can be accounted for. In particular, this one-factor model can explain the level and asymmetry of empirical exceedance correlations. However, more subtle effects require an extension of the one factor model, where the variance and skewness of the residuals also depend on the market return.

1 Introduction

Understanding the relationship between the statistics of individual stock returns and that of the corresponding index is a major issue in several finance problems such as risk management [1] or market micro-structure modeling. It is also crucial for building optimized portfolios containing both index and stocks derivatives [2, 3]. Although the index return is the (weighted) sum of stock returns, it actually displays very different statistical properties from what would result if the stock returns were independent. In particular, the cumulants (that is, the volatility, the skewness and the kurtosis) of the index distribution, which should be suppressed by a power of the number of stocks N for independent returns, are still very large, even for $N = 500$. The negative skewness of the index, in particular, is actually larger than for individual stocks, and reflects a specific leverage effect [4].

It is a common belief that cross-correlations between stocks actually *fluctuate* in time, and increase substantially in a period of high market volatility. This has been discussed in many papers – see for example [5], with more recent discussions, including new indicators, in [6, 7, 8, 9]. Furthermore, this increase is thought to be larger for large downward moves than for large upward moves. The dynamics of these correlations themselves, and their asymmetry, should be estimated, leading to rather complex models [7, 10, 11, 12]. The view of 'moving' correlations has a direct consequence for risk management: the risk for a given portfolio is seen as resulting from both volatility fluctuations and correlation fluctuations.

An alternative point of view is provided by factor models with a fixed correlation structure. The simplest version contains a unique factor – the market itself. In this case, the time fluctuations of the measured cross-correlations between stocks is, as we show below, directly related to the fluctuations of the market volatility. The notion of "correlation" risk therefore reduces to market volatility risk, which considerably simplifies the problem. In this paper, we want to address to what extent a *non-Gaussian* one-factor model is able to capture the essential features of stocks cross-correlations, in particular in extreme market conditions. Our conclusion is that most of the extreme risk correlations, measured by different indicators, are actually captured by this simple *fixed-correlation* model. This model is able to reproduce quantitatively the observed *exceedance correlations* [6] without invoking the idea of 'regime switching' recently advocated in this context in [7, 8].

However, a more detailed analysis shows that a refined model is needed to account for the dependence of the conditional volatility and skewness of the residuals on the market return.

2 A non-Gaussian one-factor model

We want to compare empirical measures of correlation with the prediction of a *fixed-correlation* model. However, for generic non-Gaussian probability distributions of returns, there is no unique way of building a multivariate process. A natural choice is to assume that the return of every stock is the sum of random independent (non-Gaussian) factors. While a multivariate Gaussian process can always be decomposed into independent factors, this is not true for generic non-Gaussian distributions. The existence of such a decomposition is thus part of the definition of our model.

The model: We will call *market* the dominant factor in this decomposition and write:

$$r_i(t) = \beta_i r_m(t) + \epsilon_i(t). \tag{1}$$

The daily return is defined as $r_i(t) = S_i(t)/S_i(t-1) - 1$, where $S_i(t)$ is the value of the stock i on day t. The return is thus decomposed into a market part $r_m(t)$ and a residual part $\epsilon_i(t)$. In a generic factor model, the residuals $\epsilon_i(t)$ are combinations of all the factors except the market and are therefore independent of it. The *one-factor model* corresponds to the simple case where the $\epsilon_i(t)$ are also independent of one another.

The market is defined as a weighted sum of the returns of all stocks. The weights can be those of a market index such as the S&P 500. These could also be the components of the eigenvector with the largest eigenvalue of the stocks cross-correlation matrix [13]. We have chosen to work simply with uniform weights, leading to the following definition:

$$r_m = \frac{1}{N} \sum_{i=1}^{N} r_i. \tag{2}$$

Had we chosen another weighting scheme for the definition of the market, the theoretical results below would still hold exactly provided that we replace averages over all stocks by the corresponding weighted averages. On our data set, the different weighted averages give essentially the same results. The coefficients β_i's are then given by:

$$\beta_i = \frac{\langle r_i r_m \rangle - \langle r_i \rangle \langle r_m \rangle}{\langle r_m^2 \rangle - \langle r_m \rangle^2}, \tag{3}$$

where the brackets refer to time averages. This model is meaningful in the case where β_i is constant or slowly varying in time. Eq. (2) immediately implies $(1/N) \sum_{i=1}^{N} \beta_i = 1$.

An important qualitative assumption of this model is that although the market is built from the fluctuations of the stocks, it is a more fundamental quantity than the stocks themselves. Hence, one cannot expect to explain the statistical properties of the market from those of the stocks within this model.

Real data and surrogate data: The data set we considered is composed of the daily returns of 450 U.S. equities among the most liquid ones from 1993 up to 1999. In order to test the validity of a one-factor model, we also generated surrogate data compatible with this model. Very importantly, the one-factor model we consider is not based on Gaussian distributions, but rather on fat-tailed distributions that match the empirical observations for both the market and the stocks daily returns [14].

The procedure we used to generate the surrogate data is the following:

- Compute the β_i's using Eq. (3) over the whole time period [15]. These β_i's are found to be rather narrowly distributed around 1, with $(1/N)\sum_{i=1}^{N}\beta_i^2 = 1.05$.
- Compute the variance of the residuals $\sigma_{\epsilon_i}^2 = \sigma_i^2 - \beta_i^2\sigma_m^2$. On the dataset we used σ_m was 0.91% (per day) whereas the rms σ_{ϵ_i} was 1.66%.
- Generate the residual $\epsilon_i(t) = \sigma_{\epsilon_i} u_i(t)$, where the $u_i(t)$ are independent random variables of unit variance with a leptokurtic (fat tailed) distribution — we have chosen here a Student distribution with an exponent $\mu = 4$:
$$P(u) = \frac{3}{(2+u^2)^{5/2}}, \tag{4}$$
which is known to represent adequately the empirical data [14].
- Compute the surrogate return as $r_i^{\text{surr}}(t) = \beta_i r_m(t) + \epsilon_i(t)$, where $r_m(t)$ is the *true* market return at day t.

Therefore, within this method, both the empirical and surrogate returns are based on the very same realization of the market statistics. This allows us to compare meaningfully the results of the surrogate model with real data, without further averaging. It also short-cuts the precise parameterization of the distribution of market returns, in particular its correct negative skewness, which turns out to be crucial.

3 Conditioning on large returns

Conditioning on absolute market return: We have first studied a measure of correlations between stocks conditioned on an extreme market return. It is indeed commonly believed that cross-correlations between stocks increase in such "high-volatility" periods. A natural measure is given by the following coefficient:
$$\rho_>(\lambda) = \frac{\frac{1}{N^2}\sum_{i,j}\left(\langle r_i r_j\rangle_{>\lambda} - \langle r_i\rangle_{>\lambda}\langle r_j\rangle_{>\lambda}\right)}{\frac{1}{N}\sum_i\left(\langle r_i^2\rangle_{>\lambda} - \langle r_i\rangle_{>\lambda}^2\right)}, \tag{5}$$
where the subscript $>\lambda$ indicates that the averaging is restricted to market returns r_m in absolute value larger than λ. For $\lambda = 0$ the conditioning disappears. Note that the quantity $\rho_>$ is the average covariance divided by the average variance, and therefore differs from the average correlation coefficient. We have studied the latter quantity, and the following conclusions remain valid in this case also.

In a first approximation, the distribution of individual stocks returns can be taken to be symmetrical, leading $\langle r_i\rangle_{>\lambda} \simeq 0$. The above equation can therefore be transformed into:
$$\rho_>(\lambda) \simeq \frac{\sigma_m^2(\lambda)}{\frac{1}{N}\sum_{i=1}^{N}\sigma_i^2(\lambda)}, \tag{6}$$

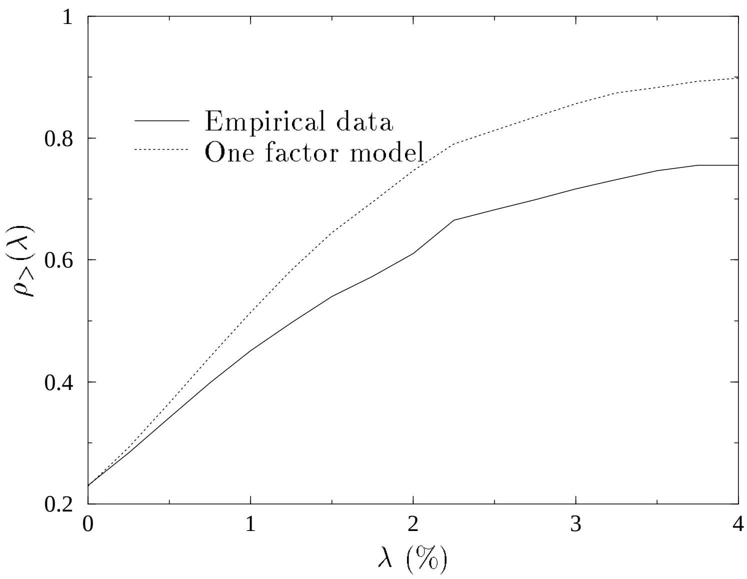

Figure 1: Correlation measure $\rho_>(\lambda)$ conditional to the absolute market return to be larger than λ, both for the empirical data and the one-factor model. Note that both show a similar apparent increase of correlations with λ. This effect is actually overestimated by the one-factor model with fixed residual volatilities. λ is in percents.

where $\sigma_m^2(\lambda)$ is the market volatility conditioned to market returns in absolute value larger than λ, and $\sigma_i^2(\lambda) = \langle r_i^2 \rangle_{>\lambda} - \langle r_i \rangle_{>\lambda}^2$. In the context of a one-factor model, we therefore obtain:

$$\rho_>(\lambda) = \frac{\sigma_m^2(\lambda)}{\left(\frac{1}{N}\sum_{i=1}^N \beta_i^2\right)\sigma_m^2(\lambda) + \frac{1}{N}\sum_{i=1}^N \sigma_{\epsilon_i}^2}. \tag{7}$$

The residual volatilities $\sigma_{\epsilon_i}^2$ are independent of r_m and therefore of λ whereas $\sigma_m^2(\lambda)$ is obviously an increasing function of λ. Hence the coefficient $\rho_>(\lambda)$ is an increasing function of λ. The one-factor model therefore predicts an increase of the correlations (as measured by $\rho_>(\lambda)$) in high volatility periods. This conclusion is quite general, it holds in particular for any factor model, even with Gaussian statistics. Therefore, the very fact of conditioning the correlation on large market returns leads to an increase of the measured correlation. A similar discussion in the context of Gaussian models can be found in [6].

More precisely, we can now compare the coefficient $\rho_>(\lambda)$ measured empirically to one obtained within the one-factor model defined above. This is presented in Fig. 1. Interestingly, the surrogate and empirical correlations are similar, displaying qualitatively the same increase of the cross-correlation when conditioned to large market returns. This shows that a one-factor model does indeed account quantitatively for the apparent increase of cross-correlations in high volatility

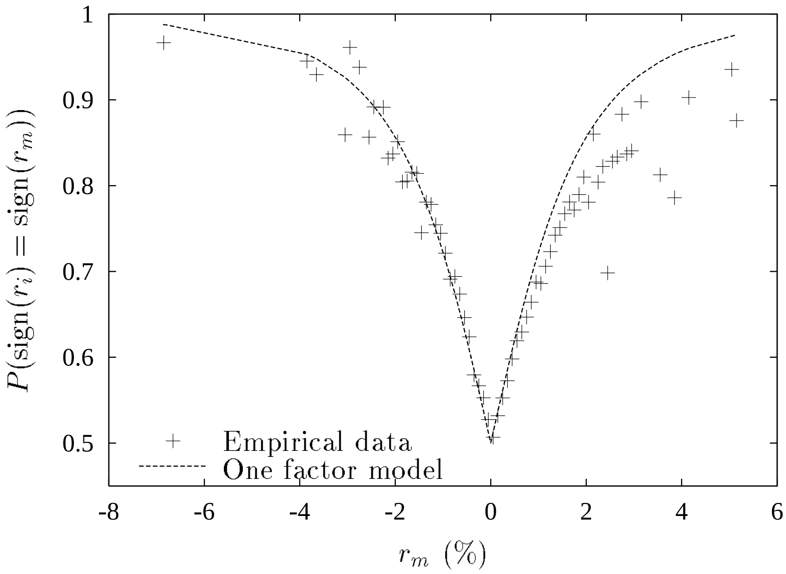

Figure 2: Conditional probability that a stock has the same sign as the market return as a function of the market return; r_m is in percent. Each cross represents the empirical probability using 4% of the days centered around a given market return. The dotted line is the prediction of the non-Gaussian one-factor model.

periods.

The one-factor model actually even *overestimates* the correlations for large λ. This overestimation can be understood qualitatively as a result of a positive correlation between the amplitude of the market return $|r_m|$ and the residual volatilities σ_{ϵ_i}, which we discuss in more details in Section 4 below (see in particular Fig. 5). For large values of λ, σ_{ϵ_i} is found to be larger than its average value. From Eq. (7), this lowers the correlation $\rho_>(\lambda)$ as compared to the simplest one-factor model where the volatility fluctuations of the residuals are neglected.

Conditional fraction of positive/negative returns: Another quantity of interest is the fraction of stocks returns having the same sign as the market return, as a function of the market return itself. The empirical results are shown on Fig. 2. We observe that for the largest returns, 90% of the stocks have the same return sign as that of the market. Therefore, the sign of the market appears to have a very strong influence on the sign of individual stock returns.

This fraction can be calculated exactly within the one-factor model. Focusing on positive market return (the case of negative returns can be treated similarly), a stock return r_i is positive whenever $\epsilon_i > -\beta_i r_m$. Therefore the average fraction $f(t)$ of stocks having a positive return for a given market return $r_m(t)$ is

$$f(t) = \frac{1}{N} \sum_{i=1}^{N} \mathcal{P}_< \left(\frac{\beta_i r_m(t)}{\sigma_{\epsilon_i}} \right), \tag{8}$$

where $\mathcal{P}_<$ is the cumulative normalized distribution of the residual (chosen here to be a Student distribution with a exponent $\mu = 4$). $f(t)$ is also plotted on Fig. 2 and fits well the empirical results. The theoretical estimate slightly overestimates the fraction $f(t)$ for positive market returns. As explained above, the correlations between σ_{ϵ_i} and $|r_m|$ do lower $f(t)$ as needed for the positive side. However, the corresponding fraction for the negative side would then be underestimated.

Conditioning on large individual stock returns – quantile correlations and exceedance correlations: Since the volatility of the residuals is two times larger than the volatility of the market, the conditioning by extreme market events does not necessarily select extreme individual stock moves. The quantities studied in the previous section, namely return correlations and sign correlations, are therefore more related to the central part of the stocks distribution rather than to their extreme tails. We now study more specifically how extreme stock returns are correlated between themselves. A first possibility is to study *quantile correlations*, that we define as:

$$\rho(q) = \frac{\frac{1}{N^2}\sum_{i,j}\left(\langle r_i r_j\rangle_q - \langle r_i\rangle_q\langle r_j\rangle_q\right)}{\frac{1}{N}\sum_i\left(\langle r_i^2\rangle_q - \langle r_i\rangle_q^2\right)}, \tag{9}$$

where the subscript q indicates that we only retain in the average days such that *both* $|r_i|$ and $|r_j|$ take their $q-$quantile value, within a certain tolerance level (this tolerance is taken to be 4% of the total interval for each quantile). In the limit $q \to 1$, this selects extremes days for both stocks i and j simultaneously. The empirical results for $\rho(q)$ are compared with the prediction of the one-factor model in Fig. 3. The agreement is again very good, though the one-factor model still slightly overestimates the true correlations in the extremes.

Another interesting quantity that has been much studied in the econometric literature recently, is the so-called exceedance correlation function, introduced in [6]. One first defines normalized centered returns $\tilde{r}_i$ with zero mean and unit variance. The positive exceedance correlation between i and j is defined as:

$$\rho_{ij}^{+}(\theta) = \frac{\langle \tilde{r}_i\tilde{r}_j\rangle_{>\theta} - \langle \tilde{r}_i\rangle_{>\theta}\langle \tilde{r}_j\rangle_{>\theta}}{\sqrt{(\langle \tilde{r}_i^2\rangle_{>\theta} - \langle \tilde{r}_i\rangle_{>\theta}^2)(\langle \tilde{r}_j^2\rangle_{>\theta} - \langle \tilde{r}_j\rangle_{>\theta}^2)}}, \tag{10}$$

where the subscript $> \theta$ means that both normalized returns are larger than θ. Large θ's correspond to extreme correlations. The negative exceedance correlation $\rho_{ij}^{-}(-\theta)$ is defined similarly, the conditioning being now on returns smaller than $-\theta$. Fig. 4 shows the exceedance correlation function, averaged over the pairs i and j, both for real data and for the surrogate one-factor model data. As in previous papers, we have shown $\rho_{ij}^{+}(\theta)$ for positive θ and $\rho_{ij}^{-}(-\theta)$ for negative θ. As in previous studies [6, 7, 8], we find that $\rho^{\pm}(\pm\theta)$ *grows* with $|\theta|$ and is larger for large negative moves than for large positive moves. This is in strong contrast

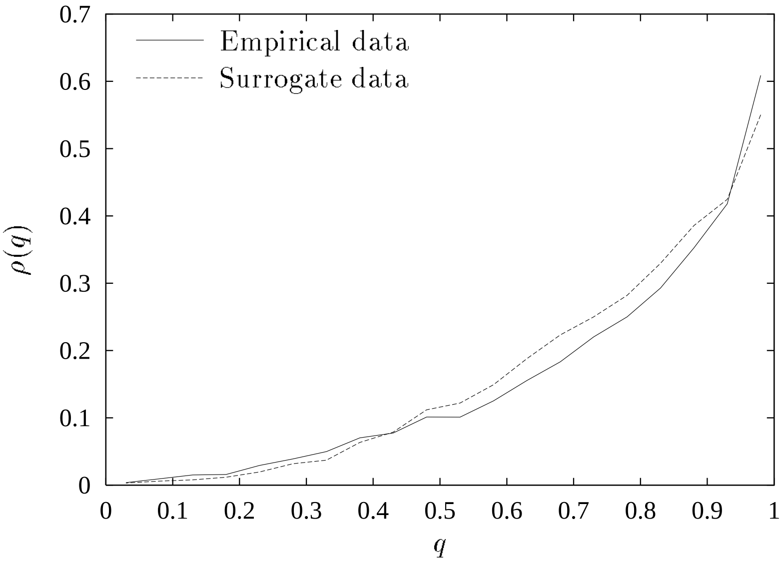

Figure 3: Correlation between stocks for joint extreme moves, $\rho(q)$, as a function of the quantile value q, both for real data and the surrogate one-factor model.

with the prediction of a Gaussian model, which gives a symmetric tent-shaped graph that goes to zero for large $|\theta|$. Note however that previous studies have focused on fixed pairs of assets i and j (for example a few pairs of international markets). The result of Fig. 4 is interesting since it reveals a systematic effect over all pairs of a pool of 450 stocks.

Several models have been considered to explain the observed results [7, 8]. Simple GARCH or Jump models cannot account for the shape of the exceedance correlations. Qualitatively similar graphs can however be reproduced within a rather sophisticated 'regime switching' model, where the two assets switch between a positive, low volatility trend with small cross-correlations and a negative, high volatility trend with large cross-correlations. Note that by construction, this 'regime switching' model induces a strong skew in the 'index' (i.e. the average between the two assets). Fig. 4 however clearly shows that a *fixed* correlation non Gaussian one-factor model is enough to explain quantitatively the level and asymmetry of the exceedance correlation function. In particular, the asymmetry is induced by the large negative skewness in the distribution of index returns, and the growth of the exceedance correlation with $|\theta|$ is related to distribution tails fatter than exponential (in our case, these tails are indeed power-laws).

4 Conditional statistics of the residuals

We conclude from the above results that the observed fluctuations of the stock cross-correlations are mainly a consequence of the volatility fluctuations and

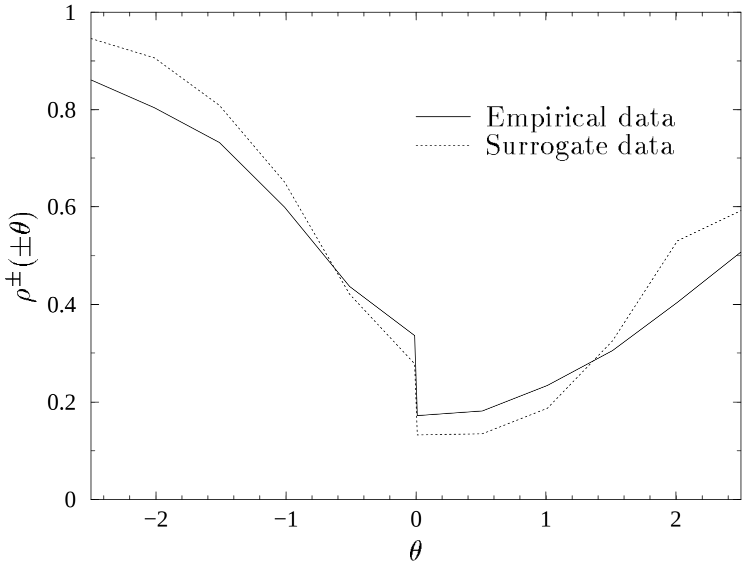

Figure 4: Average exceedance correlation functions between stocks as a function of the level parameter θ, both for real data and the surrogate one-factor model. We have shown $\rho_{ij}^{+}(\theta)$ for positive θ and $\rho_{ij}^{-}(-\theta)$ for negative θ. Note that this quantity *grows* with $|\theta|$ and is strongly asymmetric.

skewness of the market return, and that a non Gaussian one-factor model does reproduce satisfactorily most of the observed effects. However, some small systematic discrepancies appear, and call for an extension of the one-factor model. The most obvious effect not captured by a one-factor model is the recently discovered 'ensemble' skewness in the daily distribution of stock returns, as discussed by Lillo and Mantegna [16]. More precisely, they have shown that the histogram of all the stocks returns *for a given day* displays on average a positive skewness when the market return is positive, and a negative skewness when the market return is negative. The amplitude of this skewness furthermore grows with the absolute value of the market return. Note that this skewness is *not* related to the possible non zero skewness of individual stocks that has been recently discussed in several papers in relation with extended CAPM models [12, 17].

Clearly, this 'ensemble' skewness that depends on the market return cannot be explained by the above one-factor model where the residuals have a time independent zero skewness. The one-factor model is certainly an oversimplification of the reality: although the market captures the largest part of the correlation between stocks, industrial sectors are also important, as can be seen from a diagonalization of the correlation matrix [18]. Large moves of the market can be dominated by extreme moves of a single sector, while the other sectors are relatively unaffected. This effect does induce some skewness in the fixed-day histogram of stock returns distribution.

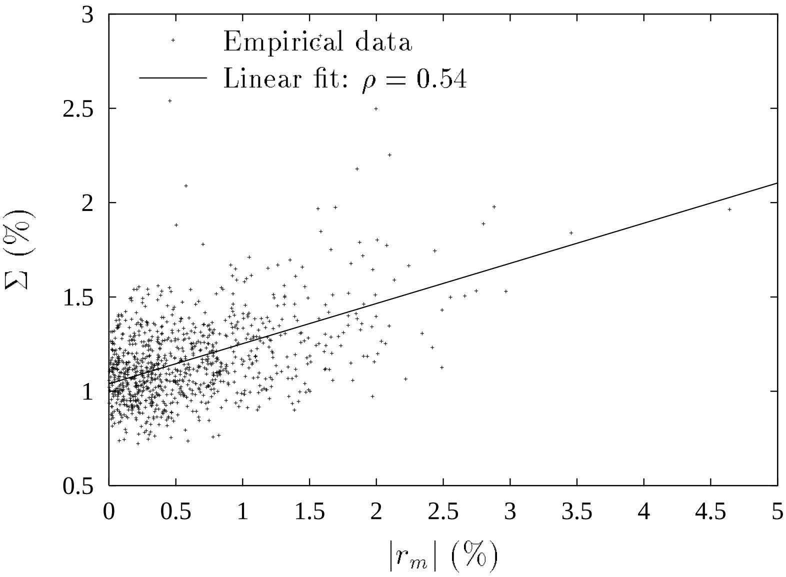

Figure 5: Daily residual 'volatility' Σ (Eq. (11)) averaged over the different stocks for a given day as a function of the market return for the same day. Σ is in percent.

A way to account for this effect is to allow the distribution of the residual $\epsilon_i(t)$ to depend on the market return $r_m(t)$. In order to test this idea, we have studied directly some moments of the distribution of the residuals *for a given day* as a function of the market return that particular day. We have studied the following quantities:

$$\Sigma = [|\epsilon_i - [\epsilon_i]|], \tag{11}$$

$$\mathcal{S} = \frac{[\epsilon_i] - \mathrm{Med}(\epsilon_i)}{\Sigma}, \tag{12}$$

$$\mathcal{K} = \frac{[\epsilon_i^2] - [\epsilon_i]^2}{\Sigma^2}, \tag{13}$$

where the square brackets [...] means that we average over the different stocks for a given day and Med selects the median value of ϵ_i. These three quantities should be thought as *robust* alternatives to the standard variance, skewness and kurtosis, which are based on higher moments of the distribution.

The quantity Σ measures the 'volatility' of the residuals and is shown in Fig. 5 as a function of $|r_m(t)|$. A linear regression is also shown for comparison. It is clear that there is a positive correlation between the market volatility and the volatility of the residuals, not captured by the simplest one-factor model. As explained above, this effect actually allows one to account quantitatively for the systematic overestimation of the observed correlations.

In order to confirm the skewness effect of Lillo and Mantegna, we have then studied the quantity $\mathcal{S}$. This quantity is positive if the distribution is positively skewed. Fig. 6 shows a scatter plot of $\mathcal{S}$ as a function of r_m [19]. Again these

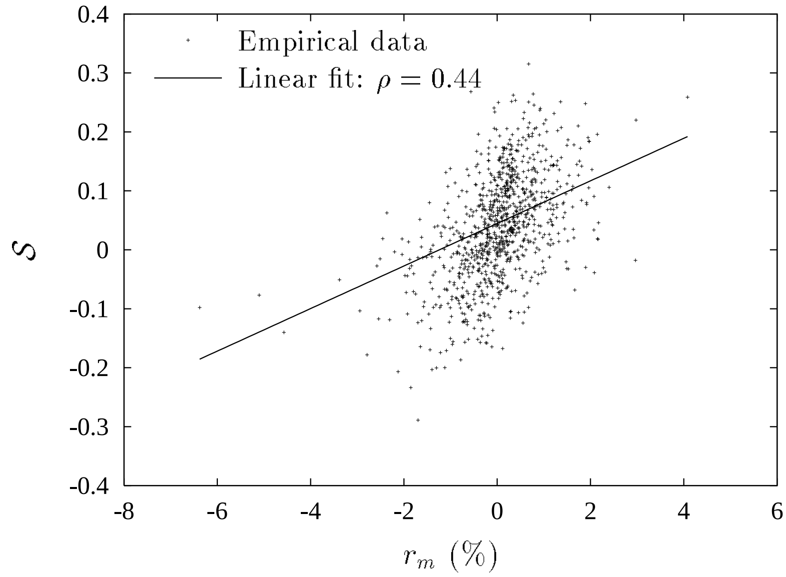

Figure 6: Daily residual 'skewness' $\mathcal{S}$ averaged over the different stocks for a given day as a function of the market return for the same day. Note that the skewness is computed using low moments of the distribution to reduce the measurement noise and does not correspond to the usual definition (see Eq. (12)).

two quantities are positively correlated, as emphasized by Lillo and Mantegna (although their analysis is different from ours).

Therefore, both the volatility and the skew of the residuals are quite strongly correlated with the market return. One could wonder if higher moments of the distribution are also sensitive to the value of r_m. We have therefore studied the quantity $\mathcal{K}$ as one possible refined measure of the shape of the distribution of residuals. This is shown in Fig. 7 and reveals a much weaker dependence than the previous two quantities.

5 Conclusion

We have thus shown that the apparent increase of correlation between stock returns in extreme conditions can be satisfactorily explained within a *static* one-factor model which accounts for fat-tail effects. In this model, conditioning on a high observed volatility naturally leads to an increase of the apparent correlations. The much discussed exceedance correlations can also be reproduced quantitatively and reflects both the non-Gaussian nature of the fluctuations and the negative skewness of the index, and *not* the fact that correlations themselves are time dependent.

This one-factor model is however only an approximation to the true correlations, and more subtle effects (such as the Lillo-Mantegna 'ensemble' skewness)

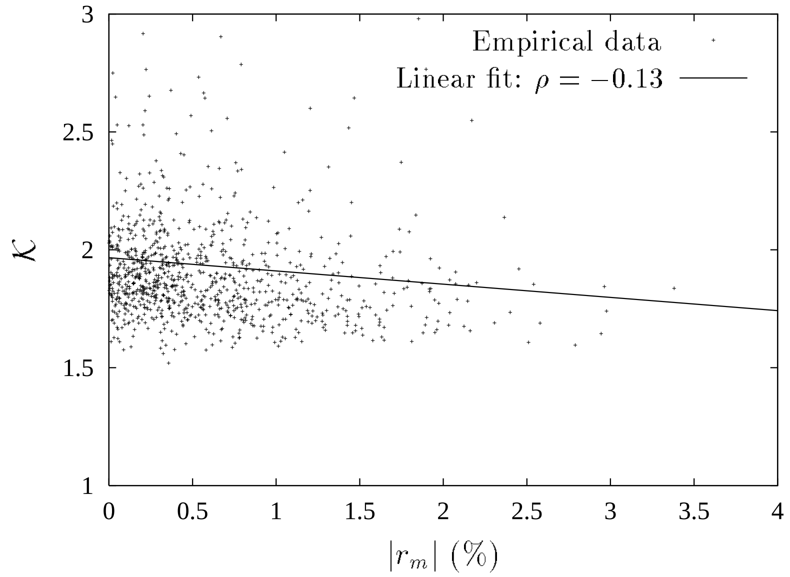

Figure 7: Daily residual 'kurtosis' averaged over the different stocks for a given day as a function of the market return for the same day. Again, the kurtosis is computed using low moments of the distribution to reduce the measurement noise and does not correspond to the usual definition (see Eq. (13)).

require an extension of the one factor model, where the variance and skewness of the residuals themselves depend on the market return.

Acknowledgments: We wish to thank M. Meyer and J. Miller for many useful discussions.

References

[1] E.J. Elton and M.J. Gruber, *Modern Portfolio Theory and Investment Analysis*, Wiley, (1995).

[2] see e.g.: J.C. Hull *Futures, Options and Other Derivatives*, Prentice Hall (2000).

[3] N. Taleb, *Dynamical Hedging*, Wiley, (1998).

[4] J.P. Bouchaud, M. Potters, A. Matacz, *The leverage effect in financial markets: retarded volatility and market panic*, e-print cond-mat/0101120.

[5] W.L. Lin, R.F. Engle, T. Ito, *Do bulls and bears move across borders? International transmission of stock returns and volatility*, in *The review of Financial Studies*, **7**, 507 (1994); C.B. Erb, C.R. Harvey, T.E. Viskanta,

Forecasting International correlations, *Financial Analysts Journal*, **50**, 32 (1994); B. Solnik, C. Boucrelle, Y. Le Fur, *International Market Correlations and Volatility*, *Financial Analysts Journal*, **52**, 17 (1996).

[6] F. Longin, B. Solnik, *Correlation structure of international equity markets during extremely volatile periods*, working paper (1999).

[7] A. Ang, G. Bekaert, *International Asset Allocation with time varying correlations*, NBER working paper (1999); A. Ang, G. Bekaert, *International Asset Allocation with Regime Shifts*, working paper (2000).

[8] A. Ang, J. Chen, *Asymmetric correlations of Equity Portfolio*, working paper (2000).

[9] This has also been argued recently by S. Drozdz, F. Grümmer, F. Ruf and J. Speth, *Dynamics of competition between collectivity and noise in the stock market*, e-print cond-mat/9911168.

[10] Y. Baba, R. Engle, D. Kraft, K. F. Kroner, *Multivariate Simultaneous Generalized ARCH*, Discussion paper 89-57, University of California, San Diego.

[11] J.Y. Campbell, A.W. Lo, A.C. McKinley, *The Econometrics of Financial Markets*, Princeton University Press (1997), and references therein.

[12] G. Bekaert, G. Wu, *Asymmetric volatility and Risk in Equity markets*, The Review of Financial Studies **13**, 1 (2000).

[13] H. Markowitz, *Portfolio Selection: Efficient Diversification of Investments* Wiley, (1959); see also: J.P. Bouchaud and M. Potters, *Théorie des risques financiers*, Aléa-Saclay, Eyrolles, (1997) (in French), *Theory of Financial Risks*, Cambridge University Press (2000); L. Laloux, P. Cizeau, J.P. Bouchaud and M. Potters, *Phys. Rev. Lett.* **83**, 1467 (1999).

[14] see e.g.: C.W.J. Granger, Z.X. Ding, *Stylized facts on the temporal distributional properties of daily data from speculative markets*, Working Paper 94-19, University of California, San Diego (1994); D.M. Guillaume *et al.*, "From the bird's eye to the microscope," *Finance and Stochastics* **1**, 2 (1997); V. Plerou, P. Gopikrishnan, L.A. Amaral, M. Meyer, H.E. Stanley, *Phys. Rev. E* **60** 6519 (1999).

[15] We do not think that computing β_i and σ_{ϵ_i} "in sample" has any consequence on the presented results.

[16] F. Lillo and R.N. Mantegna, *Symmetry alteration of ensemble return distribution in crash and rally days of financial markets*, e-print cond-mat/0002438 (2000).

[17] C. Harvey, A. Siddique, *Conditional Skewness in Asset Pricing Tests*, Journal of Finance **LV**, 1263 (2000).

[18] For a related discussion, see: R.N. Mantegna, H.E. Stanley, *An introduction to Econophysics*, Cambridge University Press (1999), Chapter 13.

[19] Due to the discreteness of quoted prices, a substantial fraction of the stocks closes at the same price two days in a row, producing a zero return. This large number of exactly zero returns often makes the median return to be zero which in turns induces an uninteresting systematic effect in the quantity $\mathcal{S}$. To get rid of this effect we have added to each return of random number uniformly distributed between -0.25% and 0.25%.