Understanding Asset Returns

L.C.G. Rogers *
Statistical Laboratory,
University of Cambridge

Liang Zhang †
Statistical Laboratory,
University of Cambridge.

September 15, 2011

First version: July 2008.
This paper appears in final form in *Mathematics and Financial Economics.*

Abstract

In this paper, we seek a model for asset returns which reproduces several well-documented stylized facts:

1. log returns are not Gaussian;
2. absolute log returns are serially correlated, but the log returns are not;
3. the Taylor effect.

There are many attempts to deal with the first, using various log-Lévy models for the asset; some of these are successful in fitting the unconditional distribution of log returns, but cannot of course reproduce the second stylized fact. We propose to model the returns with a hidden two-state Markovian regime (as in [24]), conditional on the value of which the returns have different distributions. A key observation is that if the means of the returns in the different regimes are the *same*, then the log returns are automatically uncorrelated, so we fit to index data under this restriction. By choosing symmetric hyperbolic distributions for the conditional returns, we are able to fit well the unconditional distributions, the autocovariances of absolute returns and the Taylor effect. Moreover, we find that a *common* regime model explains simultaneously these statistics for the S&P500, FTSE, DAX, Nikkei and CAC40.

Implications for investment and option pricing are discussed.

*Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 766806, email = L.C.G.Rogers@statslab.cam.ac.uk)

†Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 339798, email = L.Zhang@statslab.cam.ac.uk)

1 Introduction

The study of the asset returns has led to the settlement of some old disputes regarding the nature of the data but has also generated new challenges. A set of properties, common across many instruments, markets and time periods, has been observed by independent studies and classified as 'stylized facts'. In the light of the previous research [6, 7, 10, 13, 14, 20, 22, 26], we focus on the following properties in this paper (especially the first three).

1. **Log returns are not Gaussian:** the (unconditional) distribution of log returns seems to display a power-law or Pareto-like tail for most data sets studied.
2. **Absolute log returns are serially correlated, but the log returns are not:** autocorrelations of asset returns are often insignificant while the autocorrelation function of absolute returns decays slowly as a function of the time lag.
3. **The Taylor effect:** autocorrelations of powers of absolute return are highest at power one.
4. **Aggregational gaussianity:** as one increases the time scale over which returns are calculated, their distribution looks more and more like a normal distribution.
5. **Volatility clustering:** different measures of volatility display a positive autocorrelation indicates that high-volatility events tend to cluster in time.
6. **Gain/loss asymmetry:** large drawdowns in stock prices and stock index values but not equally large upward movements. In other words, the distribution is skewed.

The first has been studied widely and particular choices for the Lévy processes have been proposed to fit the unconditional distribution of log returns [2–4, 8, 11, 18, 19, 25]. However, the second stylized fact can not be reproduced by these methods. So we propose to model the log returns with a hidden two-state Markovian regime (HMM, also known as Regime Switching Model), conditional on the value of which the returns have different distributions.

The Regime Switching Model was first introduced by [17], followed by [16, 23, 29]. In the last two decades, this model has been used in explaining US real GNP [15], interpreting futures markets [1] and analysing foreign exchange markets [21]. In this model, it is suggested that the market returns flip between different states according to a Markov Chain; in each state, returns have the same distribution (usually Gaussian) but with different parameters. This structure can provide a simple explanation of volatility clustering.

In this paper, we propose a model with new (non-gaussian) distributions for the regimes. More importantly, we prove that if the means of the returns in the different regimes are the *same*, then the returns are automatically uncorrelated. Fitting subject to this constraint guarantees that the returns are not serially correlated; as we shall show, it proves to be possible also to match the observed serial correlation of absolute returns, which has not been achieved by the previous model [24]. The calibrated model also exhibits the Taylor effect.

Under these restrictions, we fit index data (S&P500, FTSE, DAX, Nikkei and CAC40) to various return distributions. Some of the distributions investigated are rejected by a Kolmogorov-Smirnov (KS) test; of the survivors, some do not have closed-form expressions for the characteristic function, which renders them unsuitable for option pricing. In the end, we settle on symmetric hyperbolic distributions for the conditional returns, as we find that this choice allows us to match well the stylized facts recorded above. If necessary, the regime distribtions could be extended to hyperbolic distributions to permit asymmetry of returns. Eberlein & Keller [11] prove that this distribution satisfies 'aggregational gaussianity'.

The fitting of the five indices results in estimates of the posterior probabilities of the state of the hidden Markov chain, which allows us to discuss optimal investment strategy and option pricing.

The plan of the paper is as follows. Section 2 contains the model setup, and discusses the ACF of returns and absolute returns in this model. Conditional on the state of the hidden Markov chain, the returns are drawn from some parametric distribution, and we present the candidate distributions which have been considered. The data we fitted is introduced in Section 3, followed by the calibration and statistical testing in Section 4. Section 4.2 demonstrates how well the Taylor effect is explained by fitting the model based on the symmetric hyperbolic distribution. In Section 5, we discuss various consequences of the model fitting from Section 4, in particular, how the posterior probability of the two states evolves over time; how one would optimally invest given these dynamics, and how to price options in such a model. The paper is concluded by Section 6.

2 Model Setup

2.1 Two-state Markovian Model

We choose to model the asset dynamics through a two-state hidden Markov model (HMM). There is no reason why the HMM should not have more than two states or a time-varying transition matrix. However, previous research [1,24] into three-state chains, or time-varying jump intensities fails to reveal any conclusive advantage. Set against this is the additional complexity of estimation and filtering, which can be expected to be a substantial technical obstacle. Moreover, we find that the much simpler two-state chain already does a perfectly adequate job, which is why we stick to a two-state Markovian model with constant transition matrix.

We work throughout in discrete time, as our data will be daily prices. The modelling assumption is that there is an unobserved ergodic Markov chain $(\xi_n)_{n\in\mathbb{Z}}$ which takes values in $I = \{1, 2\}$, moving between the states according to the transition matrix P. The invariant law of ξ will be denoted by π. Independently of ξ, we have two sequences $(X^i_n)_{n\in\mathbb{Z}}$ of independent random variables, $i = 1, 2$, with $X^i_n \sim F_i$ for all n and i, in terms of which the return r_n on day n is expressed as

$$r_n = \sum_{i=1}^{2} \mathbf{I}_{\{\xi_n = i\}} X_n^i. \tag{2.1}$$

We shall shortly discuss various forms for the distribution functions F_i; for now, we shall write

$$\mu_i = \int x \; F_i(dx)$$

for the mean of the distribution F_i.

According to Granger *et al.* [14], there is no empirical evidence to reject the hypothesis that asset returns are uncorrelated. The following elementary result guarantees that if the two means μ_1 and μ_2 are equal, then the model has this property.

Proposition 1. *Suppose that $\mu_1 = \mu_2 = \mu$. Then*

$$\mathbb{E}[r_n r_{n+k}] = \mu^2 \tag{2.2}$$

for any $k > 0$ and $n \in \mathbb{Z}$.

Proof. Fix $k > 0$ and let $\mathcal{X} \equiv \sigma(\xi_m, m \in \mathbb{Z})$. From (2.1), we see that

$$\begin{aligned} \mathbb{E}[r_n r_{n+k}] &= \mathbb{E}[\mathbb{E}[r_n r_{n+k} \,|\, \mathcal{X}]] \\ &= \sum_{i=1}^{2} \sum_{j=1}^{2} \mathbb{E}[\mathbb{E}[X_n^i X_{n+k}^j \,|\, \mathcal{X}]; \xi_n = i, \xi_{n+k} = j] \\ &= \sum_{i=1}^{2} \sum_{j=1}^{2} \mathbb{E}[\mu_i \mu_j; \xi_n = i, \xi_{n+k} = j] \\ &= \mu^2, \end{aligned}$$

using the fact that the X's are independent of $\mathcal{X}$ and of each other, and then using the hypothesis that $\mu_1 = \mu_2$.

$\square$

Remarks. (i) The proof does not require that the chain has two states; any Markov chain ξ for which the conditional means μ_j are the same for all j will have uncorrelated returns.

(ii) The model of Rydén *et al* [24] follows a similar modelling path, using a two-state or three-state hidden Markov chain and conditionally independent returns driven by that chain. The main differences between our work and theirs is that they assume that the distributions F_1 and F_2 are zero-mean Gaussians, whereas we consider other distributional forms, and allow the (common) mean to be non-zero. We shall also show that a common hidden Markov chain can be used to explain the dynamics of several indices at once.

2.2 Autocovariance of Absolute Returns

The autocorrelation of absolute returns has been found to decay quite slowly with lag (Granger *et al.* [14]). In the model we propose, if we set

$$\nu_i = \int |x - \mu| \; F_i(dx) \tag{2.3}$$

for the (centered) absolute first moment in regime i, we find that $\mathbb{E}|r_n - \mu| = \pi_1\nu_1 + \pi_2\nu_2$, and

$$\mathbb{E}|(r_n - \mu)(r_{n+k} - \mu)| = \begin{pmatrix} \pi_1\nu_1 & \pi_2\nu_2 \end{pmatrix} P^k \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \tag{2.4}$$

It now follows that the covariance of the centred absolute returns is given by (for $k > 0$)

$$\begin{aligned} \operatorname{cov}(|r_n - \mu|, |r_{n+k} - \mu|) &= \begin{pmatrix} \pi_1\nu_1 & \pi_2\nu_2 \end{pmatrix} \left(P^k - \begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} \pi_1 & \pi_2 \end{pmatrix} \right) \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} && (2.5) \\ &= \begin{pmatrix} \pi_1\nu_1 & \pi_2\nu_2 \end{pmatrix} v\lambda^k u^T \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}, && (2.6) \end{aligned}$$

where λ is the eigenvalue of P different from 1, and v (respectively, u) is the right (respectively, left) eigenvector of λ.

It is an inevitable consequence of our modelling assumptions that *autocorrelations of centred absolute returns decay geometrically with lag.* This appears to be in contradiction of the findings quoted in the introduction. We shall let the calibration decide for us. Visually, a slow geometric decay and a slow polynomial decay look quite similar over a reasonably long range, so if the non-unit eigenvalue $\lambda = p_{11} + p_{22} - 1$ is quite close to 1, then we stand a chance of making a reasonable approximation to the empirical autocovariance. Rydén *et al.* [24] find poor agreement in their model (see Figure 4 in their paper), and we find we do better. The main reason for this is that our calibration objective includes a term penalizing the failure of the model to fit the observed ACF of absolute returns, since this is something that we want the calibrated model to match. The use of penalized likelihood methods is widespread, as in (for example) the Akaike Information Criterion; see for example [12] and references therein.

2.3 Conditional distributions of returns

We consider conditional distributions F_i of returns given the state of the Markov chain ξ which are members of the generalized hyperbolic class of distributions, or of some subclass. The generalized hyperbolic class (GH) is a flexible class of distributions, introduced in [5], see also [11]. The density of $\mathbf{GH}(\lambda, \alpha, \beta, \delta, \mu)$ is

$$x \mapsto \frac{(\gamma/\delta)^\lambda}{\sqrt{2\pi} K_\lambda(\delta\gamma)} \frac{K_{\lambda-1/2}\left(\alpha\sqrt{\delta^2 + (x-\mu)^2}\right)}{\left(\sqrt{\delta^2 + (x-\mu)^2}/\alpha\right)^{1/2-\lambda}} \; e^{\beta(x-\mu)} \tag{2.7}$$

where $\gamma \equiv \sqrt{\alpha^2 - \beta^2}$, and the moment-generating function (MGF) is

$$z \mapsto \frac{e^{\mu z}\,\gamma^{\lambda}}{(\alpha^2 - (\beta+z)^2)^{\lambda/2}} \frac{K_\lambda(\delta\sqrt{\alpha^2 - (\beta+z)^2})}{K_\lambda(\delta\gamma)}. \tag{2.8}$$

The tails of the generalized hyperbolic distribution are exponential, as can be seen from the well-known large-x asymptotic

$$K_\nu(x) \sim e^{-x}\,\sqrt{\frac{\pi}{2x}}; \tag{2.9}$$

see p202 of [30]. Various subfamilies of the GH class are of interest in their own right:

(i) taking $\beta = 0$ gives the symmetric generalized hyperbolic class;

(ii) taking $\lambda = 1$ gives the hyperbolic class;

(iii) taking $\lambda = 1$ and $\beta = 0$ gives the symmetric hyperbolic class;

(iv) taking $\delta = 0$ and $\beta = 0$ gives the symmetric variance-gamma class;

(v) taking $\alpha = \beta = 0$ and $\lambda = -\nu/2$ gives a Student-t_ν distribution.

Thus it can be seen that the GH family contains a wide range of possible distributional shapes. We investigated quite a number of candidate conditional distributions, including the mixture of two gaussians considered by Rydén *et al.*, and eventually settled on just four classes after eliminating classes which appeared to be unable to match the data adequately: symmetric variance-gamma, symmetric hyperbolic, symmetric generalized hyperbolic and hyperbolic. In choosing these classes of distributions, we were guided by several considerations:

1. ease of calculation of absolute moments;
2. smallest number of parameters consistent with fitting the data;
3. degree of asymmetry exhibited by the data.

For the first of these, Barndorff-Nielsen [5] gives expressions for the moments and absolute moments of the generalized hyperbolic distributions: for every $\theta > 0$ and $n \in \mathbb{N}$:

$$(i) \qquad \mathbb{E}[(X-\mu)^n] = \frac{2^{\lceil \frac{n}{2} \rceil}\,(\delta\gamma)^{\lambda}\,\delta^{2\lceil \frac{n}{2} \rceil}\beta^m}{\sqrt{\pi}K_\lambda\,(\delta\gamma)\,(\delta\alpha)^{\lambda+\lceil \frac{n}{2} \rceil}} \sum_{k=0}^{\infty} \frac{2^k(\delta\beta)^{2k}\Gamma(k+\lceil \frac{n}{2} \rceil + \frac{1}{2})}{(\delta\alpha)^k(2k+m)!} K_{\lambda+k+\lceil \frac{n}{2} \rceil}(\delta\alpha) \tag{2.10}$$

$$(ii) \qquad \mathbb{E}[|X-\mu|^\theta] = \frac{2^{\frac{\theta}{2}}\,(\delta\gamma)^{\lambda}\,\delta^{\theta}}{\sqrt{\pi}K_\lambda\,(\delta\gamma)\,(\delta\alpha)^{\lambda+\frac{\theta}{2}}} \sum_{k=0}^{\infty} \frac{2^k(\delta\beta)^{2k}\Gamma(k+\frac{\theta}{2}+\frac{1}{2})}{(\delta\alpha)^k(2k)!} K_{\lambda+k+\frac{\theta}{2}}(\delta\alpha) \tag{2.11}$$

where $m := n \bmod 2$, $\lceil \frac{n}{2} \rceil := \frac{n \bmod 2 + n}{2}$ and K_λ is Bessel function of second kind. More simply, the mean of a $\mathbf{GH}(\lambda, \alpha, \beta, \delta, \mu)$ is

$$\mu + \frac{\delta\beta K_{\lambda+1}(\delta\gamma)}{\gamma K_\lambda(\delta\gamma)}. \tag{2.12}$$

For the second consideration, it is generally found that to fit univariate return distributions it is necessary to allow a four-parameter family of distributions (to set the centering, scale, right tail and left tail), and in order to fit the decay of the autocorrelation of centered absolute returns we need another two parameters, one for the rate of decay, one for the amplitude. Thus we should expect to need at least six parameters, but probably not many more, otherwise we may get overfitting. The Markov transition matrix requires two parameters to specify it (p_{11} and p_{22}), and if we had two symmetric hyperbolic distributions at our disposal, there would be a further five[1] free parameters, giving seven in total, which should be about right.

For the third consideration, it turned out that the index data that we were trying to fit did not exhibit marked asymmetry, but if we were working with individual stocks we might have found more asymmetry of returns, which might have altered our preferred choice of distribution.

3 Data Set

We apply our model to the Stock indices from the top GDP countries (S&P500, FTSE, DAX, NIKKEI, CAC40). The series consists of 5016 daily observations from 1 January 1990 to 31 December 2009. We choose 1990 as the starting date since it was the watershed for German and Japanese economics. Although China has been a major economic power for much of the 20 years analysed, there is no comparable stock index data, so we were not able to include China in the study.

Addtionally, we adjust these indices to US currency by the daily exchange rates to make the comparisons and the calibration of the model. The movements of these indices are given by Figure 1.

From Figure 1, we observe that the three European indices moved largely in step for the last two decades, and appear to rise and fall roughly in line with the US. The Japanese recession in the 1990s is visible in the rather different trajectory of the Nikkei during that time. Nevertheless, its gains and losses in the second decade fall back into line with other world markets.

The similarities of the different stock indices in Figure 1 suggest that all the markets may be driven by the similar effects, which will be modelled in our account by a *common* Markov chain. We will discuss interpretations of this common Markov chain in Section 4.

4 Calibration and results

In this section, we calibrate the parameters for different distributions by maximising the likelihood function and fitting the ACFs of absolute returns. Then we check these results with Kolmogorov-Smirnov Tests.

[1]Recall that the means are constrained to be equal.

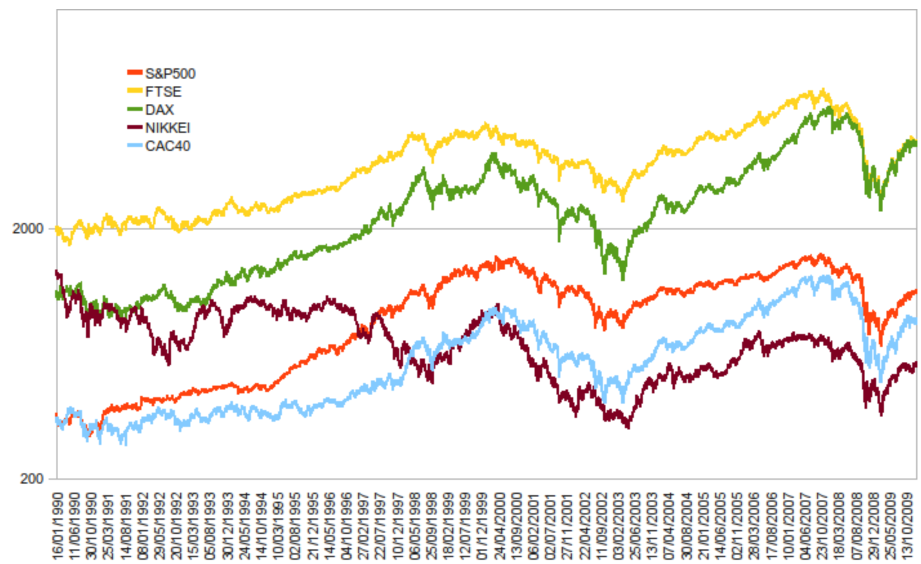

Figure 1: 1990-2009 daily Stock indices (S&P500, FTSE, DAX, NIKKEI, CAC40) adjusted by US currency

4.1 Maximum Likelihood Estimation for HMM

The log-likelihood function of an observed sequence $r_1, r_2, \ldots, r_m$ of returns is easily seen to be (see [24])

$$\mathcal{L}(\theta_1, \theta_2; r_1, \ldots, r_m) = \log\left(\pi F(r_1; \theta_1, \theta_2) P F(r_2; \theta_1, \theta_2) P \cdots P F(r_m; \theta_1, \theta_2) \mathbf{1}\right) \tag{4.1}$$

where

$$\pi = (\pi_1 \ \pi_2), \quad F(r; \theta_1, \theta_2) = \begin{pmatrix} f(r;\theta_1) & 0 \\ 0 & f(r;\theta_2) \end{pmatrix}, \quad \mathbf{1} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \tag{4.2}$$

From this, we are able to calculate maximum-likelihood estimators for the parameters, assuming that the returns are symmetric hyperbolic. Using the MLE values results in the plot shown in Figure 2 for the autocovariance of absolute returns, where the blue dashed line is the model autocovariance, and the solid black line comes from the data. It is clear that the fitted model is not matching the data well. We therefore introduce a penalty function to improve the fit:

$$\mathcal{P}(\theta_1, \theta_2) = A \sum_{k=0}^{w} (\hat{\rho}_k - \rho_k)^2 \tag{4.3}$$

where w is the total lag number for summation, A is the scale of the penalty function, and $\hat{\rho}_k$ and ρ_k are estimated and real autocovariances of absolute returns with k lags.

Explicitly, we maximize[2]

$$\mathcal{L}(\theta_1, \theta_2; r_1, \ldots, r_m) - \mathcal{P}(\theta_1, \theta_2). \tag{4.4}$$

Figure 3 shows the quality of fit achieved when we fit using $A = 10^{11}$, still assuming symmetric hyperbolic distributions for the individual conditional return distributions. How does this change when we try to fit all five indices simultaneously with the *same* driving Markov chain? The log-likelihood still has the form (4.1) where we now define

$$F(r; \theta_1, \theta_2) = \begin{pmatrix} \prod_{j=1}^{5} f(r^j; \theta_1^j) & 0 \\ 0 & \prod_{j=1}^{5} f(r^j; \theta_2^j) \end{pmatrix}, \tag{4.5}$$

and once again we maximize (4.4), where now

$$\mathcal{P}(\theta_1, \theta_2) = A \sum_{j=1}^{5} \sum_{k=0}^{w} (\hat{\rho}_k^j - \rho_k^j)^2. \tag{4.6}$$

Figure 4 presents the autocovariance of absolute returns for the S&P500 when we fit a common Markov chain to all five indices. As can be seen, the quality of fit is very similar to that achieved in Figure 3 when fitting just the S&P500 on its own.

For each of the parametric families of distributions under consideration, we firstly fitted each index on its own, then we insisted that the Markov chain was common across all indices, and fitted subject to that more exacting requirement. The results are listed in Tables 1-8. We set the penalty scalar $A = 10^{11}$ for all the cases.

The tables also report the value of the Kolmogorov-Smirnov (KS) test statistic [9] for each of the fits. To explain more fully what happened here, we firstly obtained the ML estimators of the parameters of the model, and then we calculated the supremum of the absolute differences of the empirical return distribution minus the return distribution given by the ML-fitted model. The significance levels are computed from the limiting distribution of the KS test statistic. Focusing on the KS values for the fitting of the common Markov chain, we see that all four families of distributions are doing quite well. The symmetric variance-gamma has a highest significance level of 91.46% (fitting the DAX), the symmetric hyperbolic has a highest significance level of 89.07% (again on the DAX) the symmetric generalized hyperbolic has highest significance level of 90.70% on the DAX, and the hyperbolic has highest significance level of 78.28%. There is little to choose between the first two, though the symmetric hyperbolic has a slight edge over the symmetric variance-gamma. We prefer both of these to the other two candidates, because of the smaller number of parameters that need to be fitted.

[2]It is clear that what we are doing is not a 'pure' maximum-likelihood fit. On the other hand, it is also generally well understood that MLEs are unsatisfactory for a variety of reasons, which is why penalized likelihood methods have been introduced; see, for example, [27].

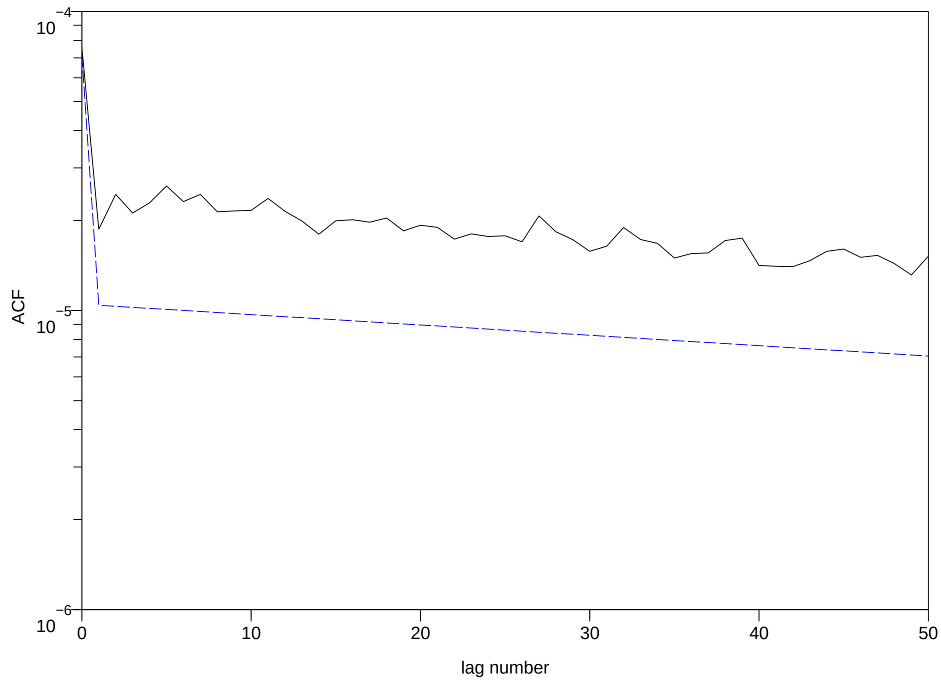

Figure 2: autocovariances of absolute return with 50 lags (1990-2009 daily S&P500)

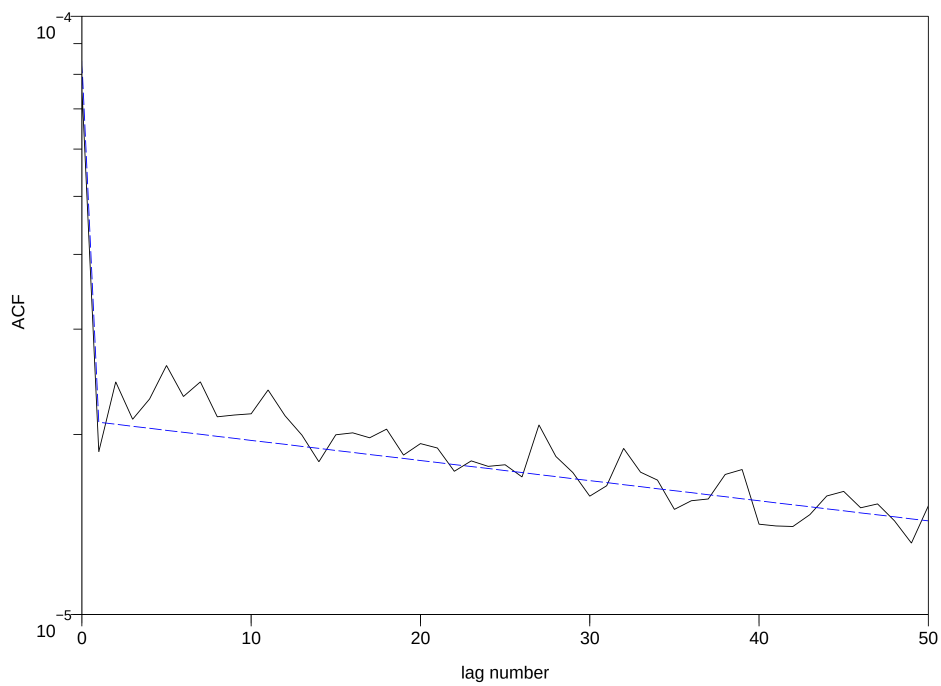

Figure 3: autocovariances of absolute return with penalty function (1990-2009 daily S&P500)

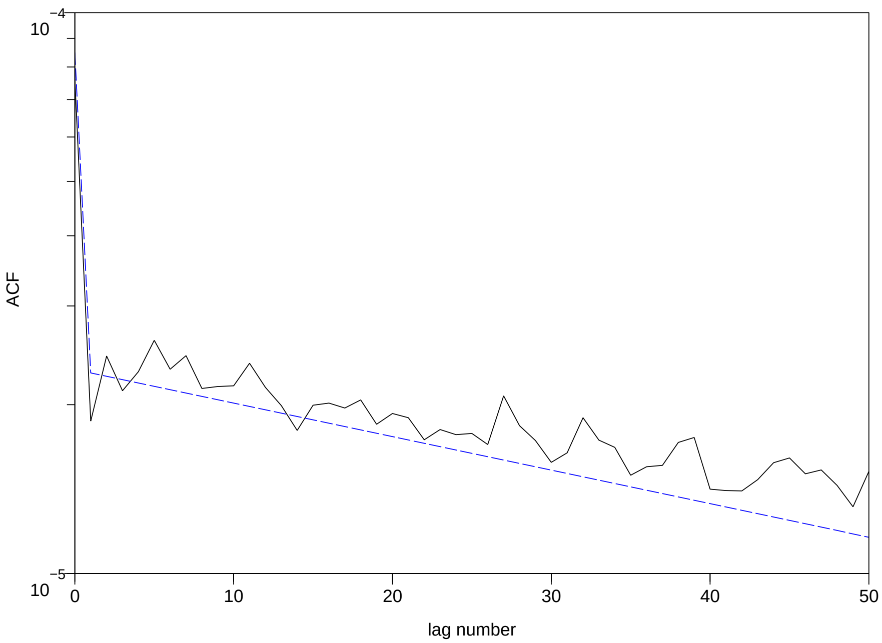

Figure 4: autocovariances of absolute returns with common Markov chain (1990-2009 daily S&P500)

4.2 Taylor Effect

A further stylized fact of returns data is the so-called 'Taylor effect', which says that the autocorrelation of powers of absolute returns are highest at power one [28]:

$$\mathbf{corr}(|r_n|, |r_{n+k}|) > \mathbf{corr}(|r_n|^\theta, |r_{n+k}|^\theta), \quad \text{for any } \theta \neq 1. \tag{4.7}$$

Here, we check this out for the absolute moments of the centred returns, using formulae (2.10) and (2.11). Plotting out the theoretical ACF[3] for different lags k and exponents θ gives us Figure 5.

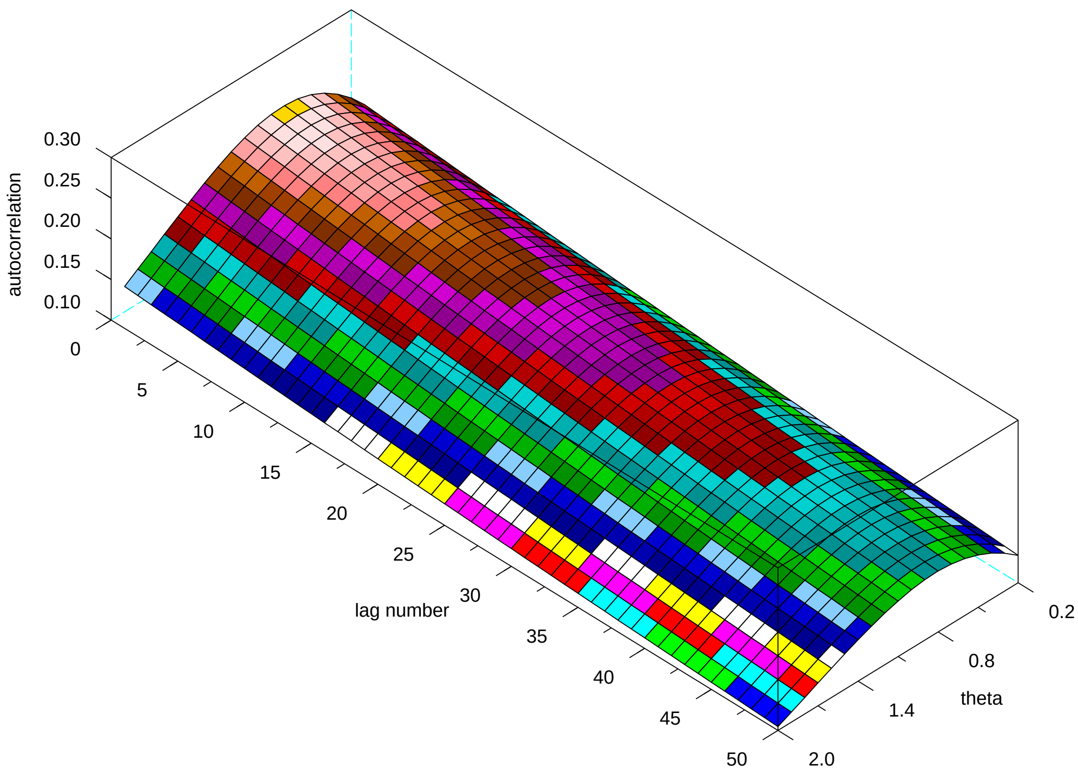

Figure 5: Autocorrelation function of $|r_n - \mu|^\theta$, $\theta = 0.2, 0.3, \ldots, 2.0$ estimated from symmetric hyperbolic distribution

As can be seen, the autocorrelation is unimodal in θ for each k, with the maximising value of θ lying in $(0.8, 1.2)$ for each k. This is consistent with the earlier findings of [24], and with the original observation of Taylor [28], which we therefore confirm within the context of our model.

[3] Here we assume the best-fitting symmetric hyperbolic model.

5 Applications

5.1 Posterior Probability

In our model, the posterior probabilities of two states are given by

$$p_n = \frac{p_{n-1} P F(r_n; \theta_1, \theta_2)}{p_{n-1} P F(r_n; \theta_1, \theta_2) \mathbf{1}} \tag{5.1}$$

where $p_n = (p_n^1, p_n^2)$ and p_n^i is the probability of being in state i at time n. Figure 6 depicts the posterior probability of being in state 2 for five indices (with the common Markov chain) from 1 January 2008 to 31 December 2009.

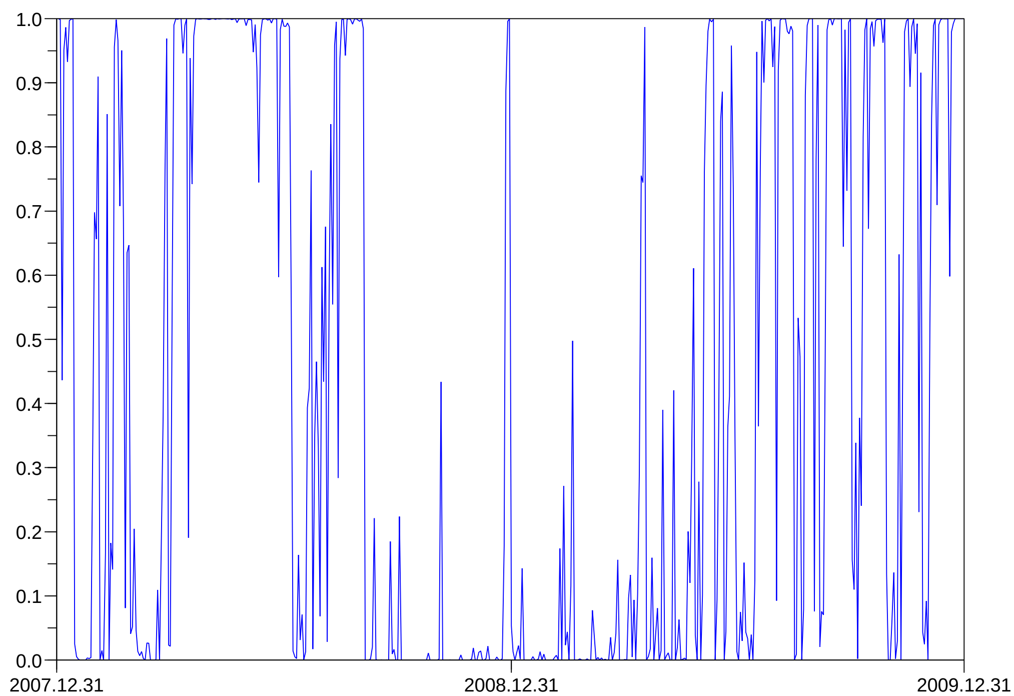

Figure 6: 2008-2009 daily S&P500 posterior probability of being optimistic.

We see that the posterior probability swings between extremely low and extremely high values, with very low values from the late summer of 2008 for almost a year. The period of very low values covers the deepest gloom of the recession, starting at about the time of the collapse of Lehman Brothers, and lasting well into the start of the Obama presidency, as more and more money was pushed into the financial system to prevent a systemic collapse. This suggests that we might interpret state 2 as an optimistic state;

it is reassuring that the hidden Markov state that we find does indeed appear to be related to major events in the world economy.

5.2 Optimal Investment

In our model, the posterior distribution of the hidden Markov state will affect investment choices; we therefore numerically solve a one-period[4] optimal investment problem. Assume an agent invests his wealth into the five stock indices (S&P500, FTSE, DAX, Nikkei and CAC40) and has CRRA utility function

$$U(x) = \frac{x^{1-R}}{1-R}, \quad R > 0. \tag{5.2}$$

He wants maximize the expectation of this utility function by choosing the portfolio π_n[5]. We assume the agent cannot short any of the indices. Then his objective function is

$$\max_{\pi_n} \mathbb{E}\left[U(w_{n+1})|p_n, w_n\right] \tag{5.3}$$

where w_n is the agent's wealth at time n, and p_n is the posterior distribution of ξ at time n. Let S_n be the vector of indicex prices at time n. Thus, w_n evolves as

$$\begin{aligned} w_{n+1} &= \pi_n^T S_{n+1} + (w_n - \pi_n^T S_n)e^r \\ &= w_n e^r + \pi_n^T S_n \left(\frac{S_{n+1}}{S_n} - e^r\right) \\ &= w_n \left(e^r + \pi_n^T \left(e^{X_{n+1}} - e^r \mathbf{1}\right)\right) \end{aligned} \tag{5.4}$$

where r is the (constant) risk-free daily return rate[6], and π_n is proportion of the agent's wealth which gets invested in the stock indices. Thus, the agent's objective becomes

$$\max_{\pi_n} \frac{w_n^{1-R}}{1-R}\, p_n P \begin{pmatrix} \int \left(e^r + \pi_n^T\left(e^{\mathbf{x}_1} - e^r\right)\right)^{1-R} f(\mathbf{x}_1;\theta_1)\, d\mathbf{x}_1 \\ \\ \int \left(e^r + \pi_n^T\left(e^{\mathbf{x}_2} - e^r\right)\right)^{1-R} f(\mathbf{x}_2;\theta_2)\, d\mathbf{x}_2 \end{pmatrix} \tag{5.5}$$

where $f(\mathbf{x}_i;\theta_i),\ i = 1,2$ are density functions of symmetric hyperbolic distribution with different parameters [7].

The numerically computed values of π are displayed in Table 9 where 'pessimistic' represents $p = (1,0)$ and 'pessimistic' stands for $p = (0,1)$ It is generally believed that the agent should cut the proportions of other assets and invest more in US ones when the market is pessimistic. Our results are consistent with this wisdom, and suggest that the agent will invest bigger proportion of the wealth in US if he is more risk-averse. Table 9 also indicates that the agent should 'put his eggs in different baskets' when the market is pessimistic.

---

[4]Because we suppose a CRRA investor, the investment decision will be the same as for a multi-period investment problem.

[5]$\pi_n = \left(\pi_n^{S\&P500}, \pi_n^{FTSE}, \pi_n^{DAX}, \pi_n^{Nikkei}, \pi_n^{CAC40}\right)$.

[6]We set the annual interest rate at 5%.

[7]Note that $\mathbf{x}_1, \mathbf{x}_2$ are vectors (with 5 rows).

5.3 Option Pricing

Now we calculate the European call option price based on our model. Since the symmetric hyperbolic distribution is not closed under convolution [11], we handle the risk-neutral distribution of multi-period returns using the characteristic function.

Assume the time to maturity of the European call option is $N \in \mathbb{N}$, the charateristic function of N-period return distribution is calculated by

$$\phi_N(t) = p_n \left(P \begin{pmatrix} \phi^1(t) & 0 \\ 0 & \phi^2(t) \end{pmatrix} \right)^N \mathbf{1} \tag{5.6}$$

where $p_n = \left(p_n^1, p_n^2\right)$ is posterior probability at time n, P is the transition matrix and $\phi^i(t)$, $i = 1, 2$ are one-period characteristic functions for two states. To convert the real probability measure into the pricing measure, we follow common practice by assuming a constant market price of risk, which simply shifts the real mean μ to the daily riskless rate r, considered as a constant. Therefore, the characteristic function of returns in the pricing measure is given by

$$\phi_N^{\mathbb{Q}}(t) = \phi_N(t) \cdot e^{itN(r-\mu)} \tag{5.7}$$

Then the risk-neutral density function is given by the Fourier inversion formula

$$f_N^{\mathbb{Q}}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \phi_N^{\mathbb{Q}}(t) dt \tag{5.8}$$

The integral on the right-hand side can be computed numerically. Here we apply FFT to get this density function and calculate the call option price by

$$C_n(N, K) = e^{-rN} \int_{-\infty}^{\infty} \left(S_n e^x - K\right)^+ f_N^{\mathbb{Q}}(x) dx \tag{5.9}$$

where K is the strike price at maturity.

6 Conclusion

This paper takes a two-state hidden Markov model for asset prices, similar to Rydén *et al.* [24], but using different families of conditional return distributions, including variance-gamma and symmetric hyperbolic, which we find work satisfactorily. Like them, we find that this simple modelling assumption is very successful at explaining many key stylized facts of asset returns, as identified by Granger *et al.* [13], [14]. However, we are able to obtain a good fit of our model to the stylized fact that the ACF of absolute returns decays quite slowly; this contrasts with the results of [24], where the empirical and fitted ACFs were substantially different.

We have fitted our model to five major stock indices simultaneously, and find that we get an excellent fit, explaining the unconditional distributions of returns in all five indices to the satisfaction of a KS test. We believe that it is important to try to do this, because if it can be made to succeed, it suggests that the rather nebulous hidden Markov state may actually be something with an economic significance; if we had to use very different Markov chains to explain each index separately, then it much harder to pin any interpretation on the states. However, we find that there *is* an economic significance to the state of the hidden Markov chain, representing pessimism or optimism. The posterior probability of state 2 during the period September 2008 to June 2009 is virtually zero, and this corresponds to the darkest months of the global financial crisis. We solve the optimal investment problem for a CRRA investor in this model, and find that in bad times there is a strong 'flight to the US' effect, again consistent with what is generally believed. Our modelling framework is also capable of pricing European options using Fourier transform techniques, and we find that here too it does a good job.

Of the six stylized facts stated in the Introduction, our modelling hypothesis deals conclusively with the first five; the last, relating to asymmetry of returns, does not arise in the index data we have been using, though if one were to work with individual stock data this would quite probably be a more prominent feature, which might require us to slightly extend the family of conditional return distributions allowed. There are many different generalizations of the standard Black-Scholes model for asset prices which have been proposed; there are, for example, GARCH models, log-Lévy models, stochastic volatility models, which all explain the stylized facts of asset returns to a greater or lesser degree. Log-Lévy models completely fail to explain the autocorrelation of absolute returns, and must be discarded. GARCH models capture many of the stylized facts, and are popular with econometricians. There are some technical issues in their use, such as the implied heavy-tailed distribution of returns, time aggregation, and the difficulty of pricing options in this modelling framework. However, their main drawback is at a conceptual level, where it is unclear how to model the feedback of returns into asset volatility in a multi-asset situation; do shocks to asset B impact on asset A, and if so, what is the mechanism? Figure 7 makes the point. Here we plot the realized quadratic variation of 29 stocks from the S&P500 index (suitably scaled to have comparable size), taking 200-day moving averages over the ten-year period July 2000 to July 2010. It is clear that the realized quadratic variation of each stock is highly variable during this ten year period, but it is also clear that periods of high volatility are common to all the stocks; *heteroskedasticity is market-wide, not stock-specific.*

The family of stochastic volatility models is also able to handle most of the stylized facts of asset returns. While such models are often formulated in a diffusion setting, what we have presented here can be considered as possibly the simplest form of stochastic volatility model - and since we impose a common mean across all the conditional return distributions, it is only the higher moments which change with the hidden state, so the stochastic volatility label is no misnomer. Our modelling framework is not difficult to work with, captures the stylized facts of asset returns, makes sense in a multi-asset situation, and offers an economically sensible interpretation of the hidden driving state. There is doubtless room for further testing, but this already seems to be a good start.

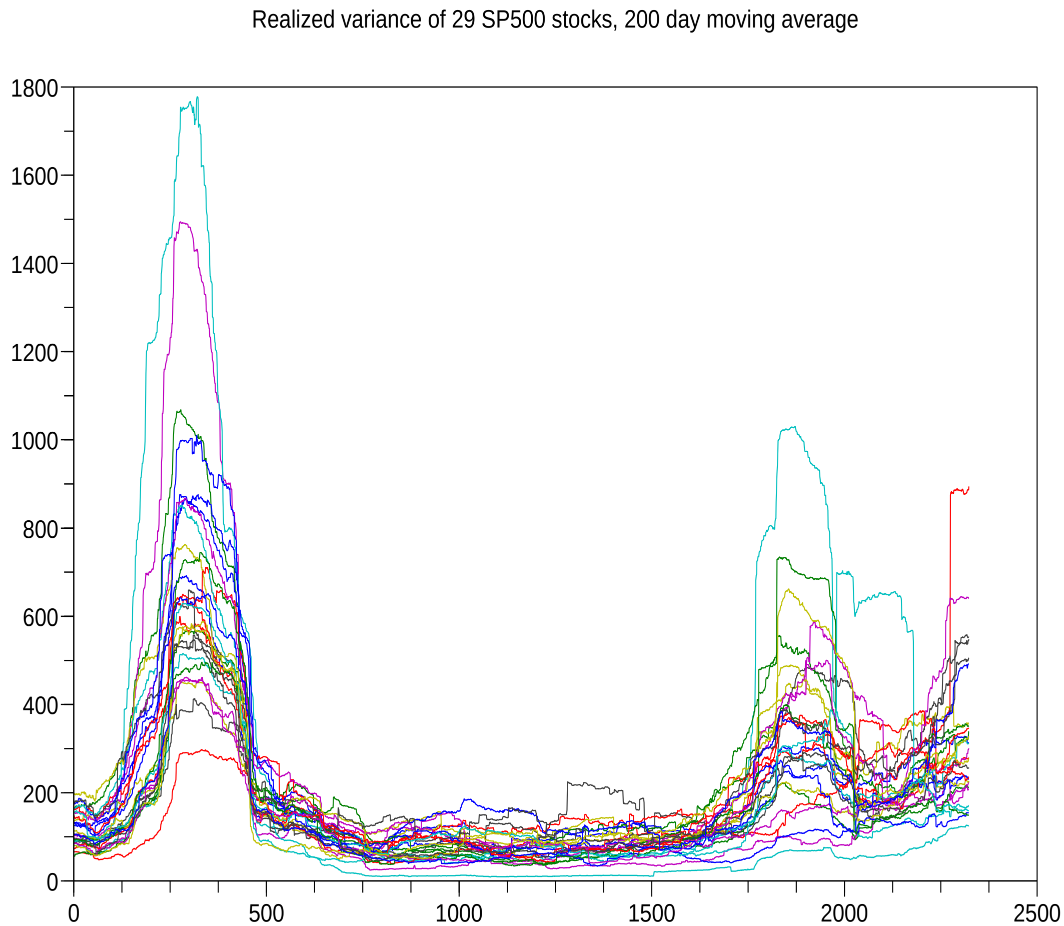

Figure 7: Heteroskedasticity.

References

[1] Alizadeh, A., and Nomikos, N. A markov regime switching approach for hedging stock indices. *The Journal of Futures Markets 24,7* (2004), 649–674.

[2] Barndorff-Nielsen, O. E. Normal inverse gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics 24,1* (1997), 1–13.

[3] Barndorff-Nielsen, O. E., Mikosch, T., and Resnick, S. *Lévy Processes-Theory and Applications*. Birkhauser, Boston, 2000.

[4] Barndorff-Nielsen, O. E., and Shephard, N. Non-gaussian ornstein-uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63,2* (2001), 167–241.

[5] Barndorff-Nielsen, O. E., and Stelzer, R. Absolute moments of generalized hyperbolic distributions and approximate scaling of normal inverse gaussian lévy processes. *Scandinavian Journal of Statistics 32,4* (2005), 617–637.

[6] Bollerslev, T., Chou, R. Y., and Kroner, K. F. Arch modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics 52,1-2* (1992), 5–59.

[7] Campbell, J., Lo, A., and Mackinlay, C. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ, 1997.

[8] Carr, P., Geman, H., Madan, D. B., and Yor, M. The fine structure of asset returns: An empirical investigation. *The Journal of Business 75,2* (2002), 305–332.

[9] Chakravarti, I. M., Laha, R. G., and Roy, J. *Handbook of Methods of Applied Statistics*. New York, Wiley, 1967.

[10] Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance 1* (2001), 223–236.

[11] Eberlein, E., and Keller, U. Hyperbolic distributions in finance. *Bernoulli 1,3* (1995), 281–299.

[12] Fan, J., and Li, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association 96* (2001), 1348–1360.

[13] Granger, C. W. J., and Ding, Z. Some properties of absolute return: An alternative measure of risk. *Annales d'Economie et de Statistique 40* (1995), 67–91.

[14] Granger, C. W. J., Spear, S., and Ding, Z. *Stylized facts on the temporal and distributional properties of absolute returns: An update*. Statistics and Finance: An Interface. Proceedings of the Hong Kong International Workshop on Statistics in Finance, 2000.

[15] Hamiltion, J. D. A new approch to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57,2* (1989), 357–384.

[16] Hamilton, J. D. Rational expectations econometric analysis of changes in regime: An investigation of the term structure of interest rates. *Journal of Economic Dynamics and Control 12* (1988), 385–423.

[17] Lindgren, G. Markov regime models for mixed distributions and switching regressions. *Scandinavian Journal of Statistics 5,2* (1978), 81–91.

[18] Madan, D. B., and Seneta, E. Chebyshev polynomial approximations and characteristic function estimation. *Journal of the Royal Statistical Society. Series B (Methodological) 49,2* (1987), 163–169.

[19] Madan, D. B., and Seneta, E. The variance gamma (v.g.) model for share market returns. *The Journal of Business 63,4* (1990), 511–524.

[20] Maddala, G., and C.Radhakrishna, R. *Handbook Of Statistics Vol 14: Statistical Methods In Finance.* Elsevier Science Ltd, Amsterdam, 1997.

[21] Marsh, I. W. High-frequency markov switching models in the foreign exchange market. *Journal of Forecasting 19* (2000), 123–134.

[22] Pagan, A. The econometrics of financial markets. *Journal of Empirical Finance 3,1* (1996), 15–102.

[23] Pagan, A. R., and Schwert, G. W. Alternative models for conditional stock volatility. *Journal of Econometrics 45,1-2* (1990), 267–290.

[24] Rydén, T., Teräsvirta, T., and Åsbrink, S. Stylized facts of daily return series and the hidden markov model. *Journal of Applied Econometrics 13* (1998), 217–244.

[25] Schoutens, W. The meixner process in finance. *EURANDOM Report 2001-002* (2001).

[26] Shephard, N. *Statistical aspects of ARCH and stochastic volatility, in D.R. Cox, David V. Hinkley and Ole E. Barndorff-Nielsen (eds.), Time Series Models in Econometrics, Finance and Other Fields, 1-67.* Chapman & Hall, London, 1996.

[27] Silverman, B. W. On the estimation of a probability density function by the maximum penalized likelihood method. *Annals of Statistics 10* (1982), 795–810.

[28] Taylor, S. *Modelling financial time series.* John Wiley & Sons: New York, 1986.

[29] Tyssedal, J. S., and Tjostheim, D. An autoregressive model with suddenly changing parameters and an application to stock market prices. *Journal of the Royal Statistical Society. Series C (Applied Statistics) 37,3* (1988), 353–369.

[30] Watson, G. N. *A Treatise on the Theory of Bessel Functions.* Cambridge University Press, Cambridge, 1966.

Table 1: Calibration for S&P500, FTSE, DAX, NIKKEI, CAC40 (symmetric variance-gamma distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.994365	0.991132	0.990237	0.982163	0.987646
p_{22}	0.997882	0.997283	0.996864	0.993746	0.997012
α_1	123.5765	126.0500	102.3876	113.8233	108.2128
α_2	288.8004	388.1931	270.9607	237.6916	328.9721
λ_1	3.203089	3.775019	3.888206	4.767314	4.478893
λ_2	2.148713	4.975458	4.204840	4.264164	6.357085
μ	0.000560	0.000436	0.000676	-0.000025	0.000450
Likelihood	16121.907	15755.240	14289.664	13680.784	14508.566
K-S test	0.800203	0.542250	0.775005	0.207451	0.200302

Table 2: Calibration with common Markov chain for S&P500, FTSE, DAX, NIKKEI and CAC40 (symmetric variance-gamma distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.989361				
p_{22}	0.997056				
α_1	104.2179	116.8230	95.5038	98.9172	98.9356
α_2	274.6175	369.5939	262.0965	194.6877	300.8480
λ_1	2.795357	3.457729	3.647810	3.764878	3.446001
λ_2	2.092718	4.470190	4.044641	3.465385	4.929850
μ	0.000500	0.000508	0.000742	0.000072	0.000520
Likelihood	74375.673				
K-S test	0.740803	0.681655	0.914569	0.756120	0.824289

Table 3: Calibration for S&P500, FTSE, DAX, NIKKEI, CAC40 (symmetric hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.994400	0.991231	0.990264	0.981943	0.987802
p_{22}	0.997879	0.997285	0.996834	0.993737	0.997180
α_1	102.5440	99.2412	82.5476	97.9610	80.6891
α_2	240.6482	297.9068	208.3344	174.8190	240.5264
δ_1	0.027138	0.030850	0.041595	0.056029	0.043254
δ_2	0.005567	0.014360	0.016282	0.017624	0.022019
μ	0.000563	0.000437	0.000681	-0.000026	0.000454
Likelihood	16121.046	15756.961	14290.552	13681.444	14509.151
K-S test	0.841059	0.527911	0.772892	0.212998	0.158088

Table 4: Calibration with common Markov chain for S&P500, FTSE, DAX, NIKKEI and CAC40 (symmetric hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.989014				
p_{22}	0.997316				
α_1	99.3756	91.1722	75.0741	76.4510	73.9903
α_2	228.9119	298.2663	207.8281	145.5122	263.6378
δ_1	0.036723	0.032127	0.043078	0.040359	0.035098
δ_2	0.006044	0.014666	0.017798	0.015606	0.024547
μ	0.000471	0.000571	0.000724	0.000025	0.000562
Likelihood	74385.057				
K-S test	0.590283	0.785565	0.890691	0.531141	0.436961

Table 5: Calibration for S&P500, FTSE, DAX, NIKKEI, CAC40 (symmetric generalized hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.994391	0.991287	0.990264	0.981976	0.987802
p_{22}	0.997866	0.997282	0.996834	0.993691	0.997140
α_1	99.6301	73.3167	82.5476	96.9700	80.6891
α_2	284.8554	296.8810	208.3344	214.4108	240.5264
δ_1	0.028314	0.042424	0.041595	0.067780	0.042883
δ_2	0.000000	0.014824	0.016282	0.000000	0.022019
λ_1	0.765809	-1.242269	0.999971	-0.180869	1.000000
λ_2	2.097148	0.842863	1.000019	3.534746	1.000000
μ	0.000561	0.000444	0.000681	-0.000035	0.000454
Likelihood	16123.433	15757.866	14290.552	13682.650	14509.151
K-S test	0.843431	0.526178	0.772859	0.126681	0.170188

Table 6: Calibration with common Markov chain for S&P500, FTSE, DAX, NIKKEI and CAC40 (symmetric generalized hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.988792				
p_{22}	0.997178				
α_1	99.3858	91.1957	75.0997	76.4506	74.0177
α_2	228.8687	298.2773	207.8440	145.3948	263.6820
δ_1	0.037271	0.031027	0.041896	0.039926	0.033512
δ_2	0.006029	0.014686	0.017504	0.015421	0.023420
λ_1	0.999993	0.999960	0.999951	0.999995	0.999995
λ_2	1.000061	0.999998	0.999997	1.000103	0.999989
μ	0.000503	0.000487	0.000726	0.000068	0.000502
Likelihood	74390.800				
K-S test	0.569885	0.527125	0.906958	0.649574	0.609072

Table 7: Calibration for S&P500, FTSE, DAX, NIKKEI, CAC40 (hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.994387	0.991204	0.990288	0.981960	0.987764
p_{22}	0.997896	0.997283	0.996843	0.993699	0.997204
α_1	101.5300	99.0992	82.1679	99.4592	80.1378
α_2	241.2241	300.0012	210.9870	174.9946	243.2188
β_1	5.400847	0.724002	1.804565	8.601457	4.315595
β_2	-8.044014	-16.743963	-13.808965	2.625868	-12.305323
δ_1	0.026507	0.030814	0.041178	0.056179	0.042716
δ_2	0.005595	0.014430	0.016483	0.017569	0.022320
μ_1	-0.001730	0.000063	-0.000708	-0.006206	-0.002936
μ_2	0.000942	0.001505	0.002196	-0.000365	0.001892
Likelihood	16122.720	15758.364	14292.825	13682.408	14510.557
K-S test	0.742916	0.277541	0.542016	0.095147	0.224072

Table 8: Calibration with common Markov chain for S&P500, FTSE, DAX, NIKKEI and CAC40 (hyperbolic distribution, penalty function for 50 lags)

	S&P500	FTSE	DAX	NIKKEI	CAC40
p_{11}	0.988756				
p_{22}	0.997117				
α_1	99.2543	97.1892	80.2113	91.1607	77.7592
α_2	234.6340	278.9116	209.1618	151.8779	212.7413
β_1	5.427253	0.030126	0.962682	6.560401	3.067956
β_2	-7.074472	-11.722466	-14.073590	-1.392728	-8.850390
δ_1	0.036583	0.034685	0.046883	0.053758	0.036340
δ_2	0.006467	0.013122	0.017625	0.017502	0.016540
μ_1	-0.002418	0.000443	-0.000148	-0.005033	-0.001778
μ_2	0.000867	0.001250	0.002351	0.000345	0.001469
Likelihood	74387.887				
K-S test	0.410540	0.459051	0.683431	0.782770	0.691301

Table 9: Investment strategies for five indices

R	2		5	
posterior probability	optimistic	pessimistic	optimistic	pessimistic
S&P500	0	0.157823	0.010009	0.203158
FTSE	0.097100	0.253471	0.316011	0.237559
DAX	0.781689	0.342397	0.464363	0.223542
Nikkei	0	0	0	0.056037
CAC40	0.121211	0.246309	0.209617	0.195069