

Stylized Facts and Simulating Long Range Financial Data¹

Laurie Davies

Fakultät Mathematik, Universität Duisburg-Essen, Germany

e-mail: laurie.davies@uni-due.de

and

Walter Krämer

Fakultät Statistik, Technische Universität Dortmund, Germany

Phone: 0231/755-3125, Fax: 0231/755-5284

e-mail: walterk@statistik.tu-dortmund.de

Abstract

We propose a new method (implemented in an R-program) to simulate long-range daily stock-price data. The program reproduces various stylized facts much better than various parametric models from the extended GARCH-family. In particular, the empirically observed changes in unconditional variance are truthfully mirrored in the simulated data.

1 Introduction and motivation

There is considerable interest in empirical finance in generating daily stock price data which mimic actual stock price behaviour as closely as possible. Such artificial data are useful, for instance, in backtesting models for value at risk or in evaluating trading strategies. The form of mimicking we shall be interested is

¹The research has been supported by the Collaborative Research Center “Statistical modeling of nonlinear dynamic processes” (SFB 823) of the German Research Foundation, which is gratefully acknowledged.

the the ability of the model to reproduce certain stylized facts about financial assets in a quantitative sense. The concept of stylized facts was introduced in [Kaldor, 1957]. There have been several papers on the application of the concept to financial data; [Rydén and Teräsvirta, 1998], [Cont, 2001], [Hommes, 2002], [Lux and Schornstein, 2005], [Bulla and Bulla, 2006], [Malmsten and Teräsvirta, 2010], [Teräsvirta, 2011]. These papers are all dynamic in that they can be used for simulations once the parameters have been estimated. In general this will require a small number of parameters as models with a large number of parameters run into estimation problems. An approach involving some form of nonparametric estimation cannot be used for simulations unless the nonparametric component can be adequately randomized. This is the approach to be taken below. The paper builds on Davies et al. (2012), who consider daily Standard and Poor's (S+P) 500 returns over 80 years. The squared returns were approximated by a piecewise constant function. This can be regarded as a nonparametric approach but in this paper we model a finer version of the piecewise constant function as a stochastic process which can then be used to simulate data.

Our main running example is the Standard and Poor's (S+P) 500 shown in the upper panel of Figure 1. The data consist of 22381 daily S+P returns with the zeros removed. The final day is 24th July 2015. The second running example is the German DAX index from 30th September 1959 to 19th October 2015 shown in the lower panel of Figure 1. There are 14049 observation of which 14026 are non-zero. A third set of data sets we shall use are the 30 firms represented in the German DAX index. The returns are from 1st January 1973, or from the date the firm was first included in the index, to 13th July 2015.

The question as to whether a model satisfactorily reproduces a quantified stylized fact or indeed any other quantified property of the data is typically answered by comparing the empirical value of a statistic with its value under the model. This was done in [Stărică, 2003] for the unconditional variance using the S+P 500 from March 4, 1957 to October 9, 2003 excluding the week starting October 19, 1987. The conclusion was that the GARCH(1,1) unconditional variance was larger than the empirical variance. For the Standard and Poor's (S+P) 500 data at our disposal the unconditional variance is 0.000135 after eliminating zero values. The maximum likelihood estimates of a GARCH(1,1)




Figure 1: Upper panel: The daily returns of S+P 500 with the zeros removed. Lower panel: the same for DAX.

model are

$$\hat{\alpha}_0 = 8.32e-07, \hat{\alpha}_1 = 0.9106, \hat{\beta}_1 = 0.08543$$

so that the unconditional second moment under the model is

$$\hat{\sigma}^2 = \hat{\alpha}_0 / (1 - \hat{\alpha}_1 - \hat{\beta}_1) = 0.000207$$

which is ‘considerably’ larger. This however ignores the variability of the second moment in simulations. On the basis of 1000 simulation the 0.05 and 0.95 quantiles of the second moment under the model are 0.000130 and 0.000345 respectively. The empirical value lies between and has an estimated p -value of 0.079 which, while small, would not be classified as statistically significant.

The same applies to the autocorrelation function. The upper panel of Figure 2 shows the ACF for the first 1500 lags for the absolute S+P 500 values in black: the grey line shows the mean of the 1000 simulations for the GARCH(1,1) model with maximum likelihood parameters. The lower panel shows nine of the 100 simulations. The large variability of the ACF values implies that comparing the empirical values with the means of simulated values can be misleading.






Figure 2: Upper panel: the first 1500 values of the empirical ACF of the absolute values of the S+P 500 data and the mean ACF of 10 GARCH simulations based on the maximum likelihood estimator for the S+P 500. Lower panel: the ACFs of the ten simulations

In [Stărică, 2003] the quantitative comparisons mentioned above were augmented by visual ones. The author compared 24 data sets generated under the model with the real data (Figure 5.1 of [Stărică, 2003]) and stated ‘The aspect of the real data is different from that of the simulated samples’. In this spirit Panel (a) of Figure 3 shows data simulated under the GARCH(1,1) model using the maximum likelihood parameters based on the whole S+P 500 data set. Panel (b) shows a simulation for the first 2000 values based on the maximum likelihood estimators for these values. The simulated data sets can be compared with the real data shown in Figure 1. The discrepancy visible in Panel (b) is very large: the average squared return is 0.398 against 0.000197 for the S+P 500 data. This is due to the fact that the maximum likelihood estimates of $\hat{\alpha}_1$ and $\hat{\beta}_1$ in the GARCH(1,1) model sum to 1.0069 so that the model is not stationary.

Such visual comparisons, also known as ‘eyeballing’, are often used (see for example [Neyman et al., 1953], [Neyman et al., 1954], [Davies, 1995], [Davies, 2008],



(a) Panel a: Simulated returns



(b) Panel b: Simulated returns for the first 2000 values

Figure 3: (a) simulated daily returns from a GARCH(1,1) model using the maximum likelihood parameter values based on the complete S+P 500 data set. (b) the same for the first 2000 values.

[Buja et al., 2009], [Davies, 2014]). Although very useful and to be recommended they have their limitations. Where possible the observed differences should be given numerical expressions and the empirical and simulated values compared. This will be done in the context of financial series in the remainder of the paper.

Whether quantitative or qualitative comparisons are made there is one fundamental problem with the S+P 500 data sets, namely that there are no independent comparable data sets. This means that it is difficult to judge the variability of such data sets. As an example some of the autocorrelations functions generated by the GARCH(1,1) process shown in Figure 2 may be judged as being too extreme to be credible for long range financial data. Figure 4 shows the first 1500 lags for the first half of the data points (lines) and the same for the second half (*). This suggests that the variability of the autocorrelation functions for the absolute returns can indeed be quite large even for very long

data sets.

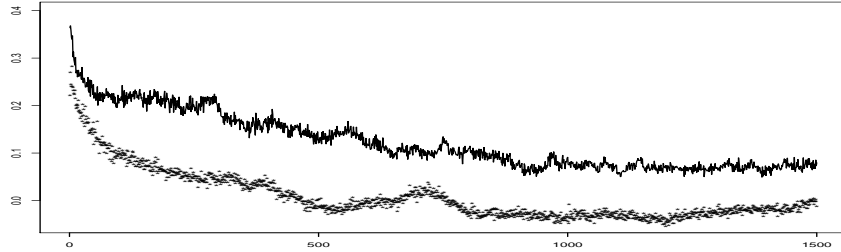


Figure 4: The first 1500 values of the empirical ACF of the S+P(500) data for the first (lines) and second (*) halves of the data.

2 Stylized facts and their quantification

In the context of financial data a list of eleven stylized facts is given in [Cont, 2001]. The ones to be considered in this paper are **1. Absence of autocorrelations**, **2. Heavy tails**, **3. Gain/loss asymmetry**, **6. Volatility clustering**, **7. Conditional heavy tails**, **8. Slow decay of autocorrelation in absolute terms** and **10. Leverage effect**. These stylized facts are exhibited by most to all of the data sets we consider. As stated above we shall be concerned with the ability of a model to reproduce these stylized facts in a quantitative sense for a given empirical time series. In some cases the quantification is straightforward, in other cases, and in particular for volatility clustering, there is no obvious manner in which this stylized fact can be quantified.

2.1 Absence of autocorrelations

The autocorrelations are not absent but small. The question is how is small to be defined. The value of the first lag for the signs of the S+P 500 data is 0.0577 which is certainly statistically significant but may not be practically relevant. The course taken in this paper is to reproduce the value of the first lag of the ACF but the software allows the user to produce other values.

Let $eacl$ denote the value of the first lag of the ACF of the signs of the data and $sac1(i), i = 1, \dots, nsim$ be the simulated values. The p -value of $eacl$

is defined by

$$p = \min(p_1, p_2) \tag{1}$$

where

$$p_1 = \#\{i : sac1(i) \leq eac1\}/nsim \text{ and } p_2 = \#\{i : sac1(i) \geq eac1\}/nsim. \tag{2}$$

The p -value is a measure of the extent to which the empirical values can be reproduced in the simulations. It is seen that $0 \leq p \leq 1/2$. This definition of a p -value will apply to any statistic whose value may be too small or too large. In some cases, only values which are too large are of interest and a one-sided definition will be used.

2.2 Heavy tails

A standard way of quantifying heavy tails is to use the kurtosis. The kurtosis is however extremely sensitive to outliers. The S+P 500 data give an example of this. The largest absolute value of the data is -0.229 and if this single value is removed the kurtosis drops from 21.51 to 15.37. Because of this extreme sensitivity the kurtosis will not be considered any further. Instead the following measure will be used. Denote the ordered absolute values of the returns normalized by their median by eaq and by aq the corresponding values for the normal distribution whereby the quantile values $qnorm(i/(n+1))$ are used. The measure of the heaviness of the tails is taken to be the mean of the difference $eaq - aq$. For normal data the value is close to zero. For $n = 23000$ the values for data with a t -distribution with 2 and 3 degrees of freedom the values are 0.451 and 0.226 respectively where the quantiles $qt(i/23001, \nu), \nu = 2, 3$ were used. The value for the S+P 500 is 0.316.

2.3 Gain/loss asymmetry

The top panel of Figure 5 shows the relative frequency of a positive return as a function of the absolute size of the return for the S+P 500 data. The centre panel shows the same for the DAX data and the bottom panel the same for Heidelberger Zement, the latter is based on the 9427 days where the return was not zero. The correlations are -0.480, -0.140 and 0.354 respectively. The plots

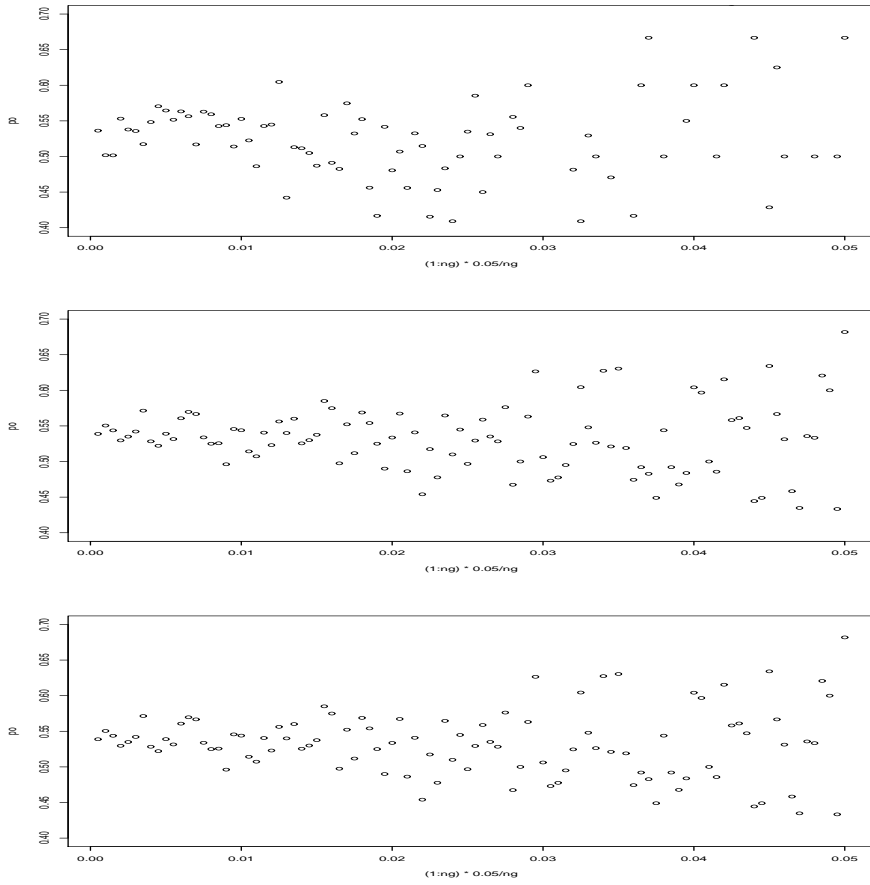


Figure 5: Top panel: proportion of positive returns as a function of the absolute return for the S+P 500 data. Centre panel: the same for the DAX data. Bottom panel: the same for Heidelberg Zement

are calculated on the basis of the 0.02-0.98 quantiles of the absolute returns. The S+P 500 and the DAX data are consistent with the remark in [Cont, 2001] that ‘one observes large drawdowns in stock prices and stock index values but not equally large upward movements’ and also ‘most measures of volatility of an asset are negatively correlated with the returns of that asset’. The Heidelberg Zement data shows that this is not always the case.

Other things being equal, which they may not be, a dependency between the absolute size of a return and its sign will induce an asymmetry in the distribution

of the returns. As a measure of symmetry we use the Kuiper distance

$$d_{\text{ku}}(\mathbb{P}_{n_+}^+, \mathbb{P}_{n_-}^-)$$

between the distributions of the positive and negative returns. This may be seen as a variant of the two-sample Kolmogorov-Smirnov test. The Kuiper values for the S+P 500, the DAX and Heidelberger Zement data sets are 0.0412, 0.0290 and 0.0342 with (asymptotic) p -values 0.000, 0.060 and 0.0810 respectively.

2.4 Volatility clustering

The quantification of volatility clustering is the most difficult stylized fact to quantify in spite of its visual clarity. The quantification we shall use is based on [Davies et al., 2012]. The basic model is

$$R_t = \Sigma_t Z_t \tag{3}$$

where Z is standard Gaussian noise. From this it follows

$$\sum_{t=i}^j \frac{R_t^2}{\Sigma_t^2} \stackrel{D}{=} \chi_{j-i+1}^2 \tag{4}$$

and hence

$$\text{qchisq}((1 - \alpha)/2, j - i + 1) \leq \sum_{t=i}^j \frac{R_t^2}{\Sigma_t^2} \leq \text{qchisq}((1 + \alpha)/2, j - i + 1) \tag{5}$$

with probability α . These latter inequalities form the basis of [Davies et al., 2012] where they are extended from one fixed interval $[i, j]$ to a family of intervals \mathcal{F} which form a local multiscale scheme. In this case the α of (5) must be replaced by α_n . The goal is to determine a piecewise constant volatility Σ_t which satisfies the inequalities (5) for all $[i, j] \in \mathcal{F}$. This problem is ill-posed. It is regularized by requiring that Σ_t minimizes the number of intervals of constancy subject to the bounds and to the values of Σ_t on an interval of constancy being the empirical volatility on that interval. Finally α_n is chosen by specifying an α and requiring that the solution is one single interval with probability α if the data are standard Gaussian white noise: see [Davies et al., 2012] for the details. For $\alpha = 0.9$ and $n = 22381$ the value of α_n is 0.9999993. For the S+P 500 data there are 76 intervals of constancy. They are shown in Figure 6.

The Z_t in (3) can be replaced by other forms of white noise, for example a t -random variable with a given number of degrees of freedom. This gives a better fit but comes at the cost of an increase in computational complexity (see Chapter 8 of [Davies, 2014]).




Figure 6: The 76 intervals of constant volatility of the the S+P 500 data.

In this paper the number of clusters will be used as a measure of the degree of clustering or the volatility of the volatility. There are other possibilities such as the sizes of the clusters (Figure 7 shows the sojourn times plotted against the volatility for the 76 intervals) but this and other measures will not be considered further.




Figure 7: Sojourn times as a function of volatility for the S+P 500.

2.5 Conditional heavy tails

The claim in [Cont, 2001] is that ‘even after correcting returns for volatility clustering the residual time series still exhibit heavy tails’. This as stated is not sufficiently precise to enable a numerical expression. If the model (3) is used with Z standard Gaussian white noise and the volatility determined as described

in Section 2.4 with $\alpha = 0.9$ then the residual times series has a kurtosis of 4.50 as against 3 for the standard normal distribution. If the Z in (3) is taken to have a t -distribution with 5 degrees of freedom and again $\alpha = 0.9$ then the residual times series has a kurtosis of 6.087 as against 6 for the t_5 -distribution. The matter will be discussed no further.

2.6 Slow decay of autocorrelation of absolute returns

As already mentioned the upper panel of Figure 2 shows the autocorrelation function of the absolute return of the S+P 500 for the first 1500 lags (black) and the mean ACF based on 100 simulations of the GARCH(1,1) model using the maximum likelihood parameters (grey). The slow decay of the ACF for the S+P 500 data is apparent.

As a measure of closeness of two autocorrelation functions a_1 and a_2 over ℓ lags we take the average absolute difference

$$d_{\text{acf}}(a_1, a_2) = \frac{1}{\ell} \sum_{i=1}^{\ell} |a_1(i) - a_2(i)|. \quad (6)$$

The p -value for the autocorrelation function of a data set based on a model is defined as follows. Let a denote the mean ACF based on the model. This can be obtained from simulations. In a second set of simulations the distribution of $d_{\text{acf}}(A, a)$ can be determined where A denotes a random ACF based on the model. Given this the p -value of $d_{\text{acf}}(ea, a)$ can be obtained relative to the distribution of $d_{\text{acf}}(A, a)$ where ea denotes the ACF of the data. For the S+P 500 data with $\ell = 1500$ $d_{\text{acf}}(ea, a) = 0.0522$ with p -value 0.061. The 0.95-quantile is 0.0571, the mean 0.0160 and the standard deviation 0.0224.

In [Teräsvirta, 2011] the value of the first lag of the ACF of the absolute returns was considered. This too will be included in the features to be reproduced.

2.7 The end return

The end return is just the end value of the stock or index given a starting value of one. We shall require that this is adequately reflected by the model. Many

models modify the basic model (3) by including an additive form for the drift

$$R_t = \mu_t + \Sigma_t Z_t \tag{7}$$

where the Z_t are assumed to have zero mean. In such a model the final return will depend on μ_t . To simulate data some stochastic assumptions must be made for μ_t and it is not clear how to do this. We prefer to keep to the basic model (3) but to let the sign of R_t depend on its absolute magnitude as in Section 2.3. It turns out that this is sufficient to successfully reproduce the end return. We point out that this can also be done for the GARCH(1,1) process without disturbing the generating scheme.

2.8 Absolute moments of the returns

Although they are not classified as stylized facts we shall require that the first and second absolute moments of the daily returns are adequately reflected by the model.

2.9 Quantiles and distribution of returns

Finally we consider two measures of the distribution of the returns. Denote by eqm , rqm and qm the order statistics of the data, of a random simulation and the mean of the simulations respectively. The mean absolute deviation of a random simulation is

$$\frac{1}{n} \sum_{i=1}^n |rqm(i) - qm(i)| \tag{8}$$

from which a one-sided p-value for the empirical deviation

$$\frac{1}{n} \sum_{i=1}^n |eqm(i) - qm(i)|$$

can be obtained.

The same applies for the Kuiper distances. With the obvious notation the simulated Kuiper distances are $d_{ku}(\mathbb{P}_{rqm}, \mathbb{P}_{qm})$ from which again a one-sided p-value can be obtained for the empirical distance $d_{ku}(\mathbb{P}_{eqm}, \mathbb{P}_{qm})$.

2.10 List of quantified features to be reproduced

In all there are eleven quantified features which are to be reproduced by the simulations. The degree to which this is accomplished will be measured by either a one-sided or a two-sided p-value as appropriate.

1. First autocorrelation of the signs of the returns
2. Heavy tails
3. Symmetry/asymmetry of returns
4. Volatility clustering - number of intervals of constant volatility
5. Slow decay of the ACF of absolute returns
6. Value of first lag of the ACF of absolute returns
7. Final return
8. Mean of absolute returns
9. Mean of squared returns
10. Quantiles of returns
11. Kuiper distance of returns

3 Modelling the data

In [Stărică, 2003] the GARCH(1,1) model is explicitly used as an example of a stationary parametric model. In the literature however it seems to be generally accepted that the S+P 500 cannot be satisfactorily modelled using this or any other stationary parametric model, see for example [Mikosch and Stărică, 2004] and [Granger and Stărică, 2005]. If this is so then alternative forms of modelling must be used. Possibilities are to use locally stationary models [Dahlhaus and Rao, 2006], segment the data and to use stationary models in each segment ([Granger and Stărică, 2005]), to use a semi-parametric approach ([David et al., 2012], [Amado and Teräsvirta, 2014]) or a non-parametric approach ([Mikosch and Stărică, 2003], [Davies et al., 2012], [David et al., 2012]) whereby the boundaries between the three approaches are

somewhat fluid. There are also more ambitious models which attempt to reproduce some stylized facts at least qualitatively by modelling the activities of the agents (see for example [Hommes, 2002], [Lux and Schornstein, 2005] and [Cont, 2007]). It seems to be difficult to adapt these to a quantitative reproduction of a particular stock.

Whether a time series is regarded as stationary, that is, it can be satisfactorily modelled by a stationary process, depends on the time horizon. Data which may not look stationary on a short horizon may be part of a data set which looks stationary on a longer horizon. Any finite data set can be embedded in a stationary process as follows. The data are extended periodically in both directions and then the origin chosen using a uniformly distributed random variable over the original data set. This may not be the best way of claiming that the data are part of a stationary process but it does show that the question of stationarity is ill-posed.

The basic model is (3). The modelling will be done in two steps, firstly modelling the volatility process Σ_t and then the ‘residual’ process Z_t .

3.1 Modelling the volatility Σ_t

The upper panel of Figure 8 shows the absolute returns of the Standard and Poor’s data together with the piecewise constant approximation of the volatility ([Davies et al., 2012]) using the value $\alpha_n = 0.9999993$. The construction of the piecewise constant volatility is a form of smoothing and as such small local variations in volatility will be subsumed in a larger interval of constant volatility. Choosing a smaller value of α_n allows the reconstruction of smaller local changes. The lower panel of Figure 8 shows the absolute returns of the Standard and Poor’s data together with the piecewise constant approximation of the volatility with $\alpha_n = 0.998$. There are 283 intervals of constancy. The choice of α_n is the first screw which can be tightened or slackened.

In a first step the log-volatilities are centred at zero by subtracting the mean. They are then approximated by a low order trigonometric polynomial

$$p_j(k) = a_j \sin(2\pi jk/n) + b_j \cos(2\pi jk/n), \quad k = 1, \dots, n \quad (9)$$

where the coefficients a_j and b_j are determined by least squares. The calcula-

tion can be made considerably faster by using the Fast Fourier Transform. The number of polynomials is determined by a further screw pow which gives the proportion of the total variance of the log-volatilities to be accounted for by the polynomial. The top panel of Figure 9 shows the logarithm of the piecewise constant volatilities centred at zero, the centre panel shows the approximating polynomial with $pow = 0.8$ which is composed of 97 polynomials of the form (9). This may be regarded as a low frequency approximation to the logarithms of the piecewise constant volatility.

The polynomials of (9) are randomized by multiplying the coefficients a_j and b_j by standard independent Gaussian random variables:

$$Z_{1j}a_j \sin(2\pi jk/n) + Z_{2j}b_j \cos(2\pi jk/n), \quad k = 1, \dots, n. \quad (10)$$

The bottom panel of Figure 9 shows a randomized version of the polynomial of the centre panel.

After removing the low frequency approximation the residuals form the remaining high frequency log-volatility. They are shown in the upper panel of Figure 10. It is not obvious how the residuals can be modelled. In the following it will be done by generating random intervals with lengths exponentially distributed with alternating means λ_1 and λ_2 . On the long intervals corresponding to λ_1 the log-volatility will be modelled as $N(0, \sigma_1^2)$ with $\sigma_1 = 0$ as the default value. On the short intervals corresponding to λ_2 the log-volatility will be modelled $\sigma_2 T_\nu$ where T_ν is t -distributed with ν degrees of freedom. The lower panel of Figure 10 shows such a randomization with $\lambda_1 = 200, \sigma_1 = 0, \lambda_2 = 20, \sigma_2 = 0.4, \nu = 15$. Adding the low and frequency components gives a randomization of the volatility as shown in Figure 11.

So far the smoothed log-volatility process has been centred at zero by subtracting the mean $mlv = -4.769$. The problem now is to specify the variability of the mean in the model. Some orientation can be obtained by dividing the data into quarters and calculating the empirical mean log-volatilities for each quarter. They are -4.451, -5.120, -4.807 and -4.694. Based on this the mean will be modelled $mlv + \Delta$ where Δ is uniformly distributed over an interval $[-\delta, \delta]$ with default value $\delta = 0.2$. This concludes the modelling of the process volatility process Δ_t .

3.2 Modelling the R_t

Let Σ_t be the volatility process described in the last section and \hat{Z}_t be i.i.d. standard normal random variables. In a first step put

$$\tilde{Z}_t = (\rho|\hat{Z}_{t-1}| + 1)\hat{Z}_t/\sqrt{1 + 2\rho\sqrt{2/\pi} + \rho^2}. \quad (11)$$

The value of ρ can be chosen so that the sixth item in the list of Section 2.10 can be adequately reproduced.

Given the \tilde{Z}_t the absolute return is set to

$$|R_t| = \Sigma_t|Z_t| = \Sigma_t|\tilde{Z}_t|(1 + |\tilde{Z}_t|)^\eta \quad (12)$$

for some value of η with default value zero. A positive value of η makes the tails of $|Z_t|$ heavier than those of the normal distribution, a negative value makes them lighter. It remains to model the sign of the return.

As is evident from Figure 5 the signs of the R_t may well depend on the value of $|R_t|$. This is taken into account as follows. Denote the $i/\nu, i = 1, \dots, \nu$ quantiles of the absolute returns of the data by $\text{qa}(i)$ and the relative frequency of the number of positive returns for those returns r_t with $\text{qa}(i-1) < |r_t| \leq \text{qa}(i)$ by $p(i)$. The default value of ν is $\nu = 50$. Given a simulated value of the absolute return $|R|$ with $\text{qa}(i-1) < |R| \leq \text{qa}(i)$ the actual return R is taken to be positive with probability $\gamma p(i) + (1 - \gamma)/2$ where $\gamma, 0 \leq \gamma \leq 1$. The parameter γ is a further screw with default value $\gamma = 1$.

Finally the first autocorrelation of the returns (first on the list of Section 2.10) can be taken into account as follows. If the empirical value of the autocorrelation is eacf1 then the final sign of R_t is determined as follows. Let $U_i, i = 2, \dots, n$ be a sequence of i.i.d. random variables uniformly distributed over $[0, 1]$. If $U_i > |\text{eacf1}|$ then the sign of R_t is unchanged. If $U_i < |\text{eacf1}|$ the the sign of R_t is set equal to that of R_{t-1} if $\text{eacf1} > 0$ and to the opposite sign of R_{t-1} if $\text{eacf1} < 0$.

4 The results of some simulations

The results for the S+P 500 and DAX data are given in Table 1. They are given in terms of the p-values for the 11 items of Section 2.10. The starred items

	Features as in Section 2.10										
	1*	2*	3*	4*	5	6*	7*	8*	9*	10	11
S+P 500	0.33	0.15	0.32	0.48	0.77	0.32	0.43	0.46	0.41	0.72	0.70
GARCH	0.32	0.06	0.17	0.12	0.05	0.30	0.48	0.00	0.07	0.03	0.00
DAX	0.33	0.16	0.24	0.46	0.88	0.50	0.48	0.29	0.30	0.54	0.51
GARCH	0.29	0.01	0.31	0.34	0.00	0.12	0.27	0.25	0.11	0.19	0.33

Table 1: The p-values (based on 1000 simulations) for the 11 items on the list of stylized facts in Section 2.10 for the modelling of Section 3 and a modified GARCH(1,1) modelling.

are two-sided p-values with a maximum value of 0.5. The GARCH(1,1) simulations have been modified to by altering the sign of the returns as described in Section 3.2. This has no effect on the absolute values of the returns and consequently no further effect on the GARCH(1,1) simulations. For the modelling described in Section 3 it proved possible in all cases to find parameter values such that all p-values exceed 0.1. No attempt was made to maximize the smallest p-values. The choice of parameter values is not easy as most of them affect several features. This problem does not occur for the GARCH(1,1) modelling. The best result for the GARCH(1,1) model is when it is applied to the DAX data. There all but two features have a p-value exceeding 0.1. The exceptions are heavy tails 2* where the GARCH(1,1) modelling results in tails which are too light. The worst failure is the inability to reproduce the slow decay of the ACF of the absolute returns, feature 5. Figure 12 shows the ACF of the DAX index (grey), the mean ACF using the modelling described in Section 3 (black) and the mean for the GARCH(1,1) modelling (dashed).

All thirty current members of the DAX were also modelled by both methods. For the modelling described in Section 3 it was always possible to choose parameter values such that all 11 features had a p-value exceeding 0.1. The GARCH(1,1) modelling turned out to be worse for these data sets than for the two indices S+P 500 and DAX. For all thirty firms the features 8-11 all had p-values of zero. In the case of 8 and 9 the model underestimated the empirical values in keep with the findings of [Stărică, 2003] for a section of the S+P 500

data. They were also underestimated for the S+P 500 data but less severely. The DAX data are exceptional in this respect, the empirical values were over-estimated.

References

- [Amado and Teräsvirta, 2014] Amado, C. and Teräsvirta, T. (2014). Modelling changes in the unconditional variance of long stock return series. *Journal of Empirical Finance*, 25:15–35.
- [Buja et al., 2009] Buja, A., Cook, D., Hofmann, H., Lawrence, M., Lee, E.-K., Swayne, D., and Wickham, H. (2009). Statistical inference for exploratory data analysis and model diagnostics. *Philosophical Transactions of the Royal Society A*, 367:4361–4383.
- [Bulla and Bulla, 2006] Bulla, J. and Bulla, I. (2006). Stylized facts of financial time series and hidden semi-markov models. *Computational Statistics and Data Analysis*, 51(4):2192–2209.
- [Cont, 2001] Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223–236.
- [Cont, 2007] Cont, R. (2007). Volatility clustering in financial markets: Empirical facts and agent-based models. In Teyssière, G. and Kirman, A. P., editors, *Long Memory in Economics*, pages 289–309. Springer.
- [Dahlhaus and Rao, 2006] Dahlhaus, R. and Rao, S. S. (2006). Statistical inference for time-varying ARCH process. *Annals of Statistics*, 34(3):1075–1114.
- [David et al., 2012] David, N., Sardy, S., and Tseng, P. (2012). ℓ_1 -penalized likelihood smoothing and segmentation of volatility processes allowing for abrupt changes. *Journal of Computational and Graphical Statistics*, 21(1):217–233.
- [Davies, 2014] Davies, L. (2014). *Data Analysis and Approximate Models*. Monographs on Statistics and Applied Probability 133. CRC Press.
- [Davies, 1995] Davies, P. L. (1995). Data features. *Statistica Neerlandica*, 49:185–245.

- [Davies, 2008] Davies, P. L. (2008). Approximating data (with discussion). *Journal of the Korean Statistical Society*, 37:191–240.
- [Davies et al., 2012] Davies, P. L., Höhenrieder, C., and Krämer, W. (2012). Recursive estimation of piecewise constant volatilities. *Computational Statistics and Data Analysis*, 56(11):3623–3631.
- [Granger and Stărică, 2005] Granger, C. and Stărică, C. (2005). Nonstationarities in stock returns. *The Review of Economics and Statistics*, 87(3):523–538.
- [Hommes, 2002] Hommes, C. H. (2002). Modeling the stylized facts in finance through simple nonlinear adaptive systems. *Proceedings of the National Academy of Sciences*, 99:7221–7228.
- [Kaldor, 1957] Kaldor, N. (1957). A model of economic growth. *The Economic Journal*, 67(268):591–624.
- [Lux and Schornstein, 2005] Lux, T. and Schornstein, S. (2005). Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics*, 41:169–196.
- [Malmsten and Teräsvirta, 2010] Malmsten, H. and Teräsvirta, T. (2010). Stylized facts of financial time series and three popular models of volatility. *European Journal of Pure and Applied Mathematics*, 3(3):443–477.
- [Mikosch and Stărică, 2003] Mikosch, T. and Stărică, C. (2003). Stock market risk-return inference: an unconditional non-parametric approach. <http://dx.doi.org/10.2139/ssrn.882820>.
- [Mikosch and Stărică, 2004] Mikosch, T. and Stărică, C. (2004). Nonstationarities in financial time series, the long-range dependence, and the IGARCH effects. *The Review of Economics and Statistics*, 86:378390.
- [Neyman et al., 1953] Neyman, J., Scott, E. L., and Shane, C. D. (1953). On the spatial distribution of galaxies a specific model. *Astrophysical Journal*, 117:92–133.
- [Neyman et al., 1954] Neyman, J., Scott, E. L., and Shane, C. D. (1954). The index of clumpiness of the distribution of images of galaxies. *Astrophysical Journal Supplement*, 8:269–294.

- [Rydén and Teräsvirta, 1998] Rydén, T. and Teräsvirta, T. and Åbsbrink, S. (1998). Stylized facts of daily return series and the hidden Markov model. *Journal of Applied Econometrics*, 13:217–244.
- [Stărică, 2003] Stărică, C. (2003). Is GARCH(1,1) as good a model as the Nobel prize accolades would imply? Technical report, Department of Mathematical statistics, Chalmers University of Technology, Gothenburg, Sweden.
- [Teräsvirta, 2011] Teräsvirta, T. and Zhao, Z. (2011). Stylized facts of return series, robust estimates and three popular models of volatility. *Applied Financial Economics*, 21(1):67–94.




Figure 8: The absolute daily returns of the Standard and Poor's index together with a piecewise constant volatility: upper panel with $\alpha_n = 0.9999993$ and 76 interval; lower panel with $\alpha_n = 0.998$ and 283 intervals.




Figure 9: Top panel: the logarithm of the piecewise constant volatilities of the lower panel of Figure 8. Centre panel: an approximating trigonometric polynomial accounting for 80% of the variance. Bottom panel: a randomized version of the polynomial.




Figure 10: Upper panel: The residuals of the log-volatility after removal of the low frequency approximation. Lower panel: a randomization of the residuals.




Figure 11: A randomization of the S+P 500 volatility process.




Figure 12: The first 1500 values of the empirical ACF of the absolute values of the DAX data (grey), the mean ACF using the modelling of Section 3 (black) and the mean ACF for the GARCH(1,1) modelling (dashed).