

Long range financial data and model choice

P. L. Davies*

Department of Mathematics, University Duisburg-Essen
Department of Mathematics, Technical University Eindhoven

February 28, 2006

Abstract

Long range financial data as typified by the daily returns of the Standard and Poor's index exhibit common features such as heavy tails, long range memory of the absolute values and clustering of periods of high and low volatility. These and other features are often referred to as stylized facts and parametric models for such data are required to reproduce them in some sense. Typically this is done by simulating some data sets under the model and demonstrating that the simulations also exhibit the stylized facts. Nevertheless when the parameters of such models are to be estimated recourse is very often taken to likelihood either in the form of maximum likelihood or Bayes. In this paper we expound a method of determining parameter values which depends solely on the ability of the model to reproduce the relevant features of the data set. We introduce a new measure of the volatility of the volatility and show how it can be combined with the distribution of the returns and the autocorrelation of the absolute returns to determine parameter values. We also give a parametric model for such data and show that it can reproduce the required features.

1 Introduction

In Section 2 we criticize the use of universal principles, in particular those based on likelihood, for determining the values of the parameters of a model given the data. We argue that the reproduction of stylized facts is a necessary but not a sufficient condition for a model to be adequate. In Section 3 a concept of approximation is introduced and it is shown how it can be used to determine parameter values. The stylized facts of Section 2 are augmented by a measure of the volatility of the volatility of the data, made more precise and incorporated into the concept of approximation. In Section 4 a stochastic model for long range financial data is introduced and it is shown how it can reproduce the chosen relevant features of the Standard and Poor's index.

*Research supported in part by Sonderforschungsbereich 475, University of Dortmund

2 Model Choice

2.1 Universal principles and likelihood

Given a data set $\mathbf{x}_n = (x_1, \dots, x_n)$ and a parametric family $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$ we call any principle for determining appropriate parameter values which is independent of the subject matter of the data a universal principle. Examples are maximum likelihood, maximum penalized likelihood, Bayes, AIC (Akaike's information criterion, Akaike (1973)), BIC (Bayes information criterion, Schwartz (1978)), MDL (Minimum description length, Rissanen (1987)), and cross validation (Stone (1974)). Apart from cross validation all these methods are likelihood based although, for the following reasons, likelihood would seem not to be an appropriate concept within the context of model choice:

- (a) Likelihood is blind.
- (b) Likelihood reduces the measure of fit between data \mathbf{x}_n and a parametric model P_θ to a single number.
- (c) It is not easy to understand on what basis likelihood determines the parameter values.
- (d) Likelihood is pathologically discontinuous with respect to the data and the model.

With respect to (a) we point out that if the data are i.i.d. Cauchy and the model is i.i.d. Gauß then the obvious inappropriateness of the model cannot be read off from the value of likelihood or the sufficient statistics. In the case of (b) the fit of any model to any data set is reduced to a single number irrespective of the complexity of the data or the model. One way of understanding (c) is to consider the Cauchy distribution. This is often used to model outliers but it is less well known that the Cauchy distribution is peaked at the origin, more so than the slash distribution which has similarly heavy tails (see for example Cohen 1991). If likelihood is used to determine the values of the parameters of the Cauchy model then it will take into account the peakedness at the origin whether this is desired or not. Finally likelihood inherits the pathological discontinuity of the differential operator. Models which are close at the level of random variables can be arbitrarily far apart in the space of densities. In other contexts likelihood has a road to play but only when it is accompanied by some form of regularization to avoid "free lunches" (see Tukey 1993). The most obvious example is the normal distribution which, in a well-defined sense, makes it most difficult to estimate the mean of a distribution. In the following section we expound an alternative approach to determining parameter values.

2.2 Long range financial data and stylized facts

Because of the aforementioned weaknesses of likelihood it is never used in its pure form when a decision about the appropriateness of a model has to be

made. Some form of diagnostics or model checking will be employed. In the case of financial data it is often required that the model reproduce certain stylized facts such as the following;

- (a) large changes in and grouping of the volatility (the top panel of Figure 1 which shows the daily returns of the Standard and Poor's index over 30 years)
- (b) heavy tails (the centre panel of Figure 1 which show the quantiles of the Standard and Poor's data plotted against the corresponding quantiles of the Gaussian distribution)
- (c) long range memory of the absolute returns (the bottom panel of Figure 1 which show the autocorrelations of the absolute returns)

Stylized facts “concentrate on broad tendencies, ignoring individual detail” (Kaldor 1961) and consequently the concept is not sufficiently precise to allow the determination of parameter values.

Their use can be illustrated by the GARCH model defined as follows. The daily returns $R(t)$ and the volatility $\Sigma(t)$ are related by

$$R(t) = \Sigma(t)Z(t) \tag{1}$$

where $Z(t)$ is standard Gaussian white noise. The model specifies

$$\Sigma(t)^2 = \alpha_0 + \alpha_1 \Sigma(t-1)^2 + \beta_1 R(t-1)^2 \tag{2}$$

where the parameters α_0 , α_1 and β_1 may be obtained from the data using for example maximum likelihood. The results for the Standard and Poor's data are

$$\alpha_0 = 0.0275, \alpha_1 = 0.9248, \beta_1 = 0.0693. \tag{3}$$

Using these parameter values it is possible to simulate a data set under the GARCH model. Corresponding to the Figure 1 we get Figure 2. It is clear that the model reproduces the stylized facts but if this is the only criterion then the same model would be appropriate for every data set which also exhibits the same stylized facts. Clearly more attention to the detail is required and we do this in the next section.

2.3 Approximating data and model choice

Following Davies (1995) we regard a stochastic model P to be an adequate approximation to a data set \mathbf{x}_n if “typical” samples $\mathbf{X}_n(P)$ of size n generated under P “look like” the sample \mathbf{x}_n . The word “typical” is specified by a number α , $0 < \alpha < 1$, with the interpretation that $100\alpha\%$ of the samples $\mathbf{X}_n(P)$ are typical. Standard values of α are the usual 0.9, 0.95 etc. The words “look like” are to be regarded as a decision rule, usually in the form of a computer programme, which decides whether a sample is typical. A simplified version of the idea is the following. Given a sample \mathbf{x}_n and a model P the statistician generates a




Figure 1: The top panel shows the daily returns of Standard and Poor's index over 30 years. The centre panel shows the quantiles of the daily returns plotted against the quantiles of the normal distribution. The bottom panel shows the autocorrelations of the absolute returns up to a lag of 300 days. The returns are standardized so that the median absolute return is 1.




Figure 2: The top panel shows a simulation of a GARCH process using the coefficients of (3). The centre panel shows the quantiles of the absolute values plotted against those of the Gaussian distribution. The bottom panel shows the autocorrelations of the absolute returns up to a lag of 300 days.

further 999 samples $\mathbf{X}_{1,n}(P), \dots, \mathbf{X}_{999,n}(P)$ giving 1000 samples in all. If $\alpha = 0.95$ the statistician specifies 950 typical samples, or equivalently, 50 atypical samples. If the real data set \mathbf{x}_n is among the typical samples the model is regarded as an adequate approximation: data generated under the model cannot be distinguished from the real data. In this context Tukey (1993) writes

...we should have to say that certain aspects of the data – not typically, but unavoidably, including “Most (modelled) observations have irrational values !” – are not to be used in relating conceptual (or simulated) samples to observed samples. Thought and debate as to just which aspects are to be denied legitimacy will be both necessary and valuable.

Those aspects which are judged to be legitimate and relevant will depend on probability considerations and on the subject matter of the data; there are no universal principles for making the decision and consequently no universal principles for model choice. An example from financial data is the share price of the German motor-car company BMW. Figure 3 shows the daily returns close to zero and a large atom at zero itself is apparent. A decision has to be made as to whether this aspect of the BMW data is to be declared legitimate or not. The data also exhibit other patterns whose cause is not immediately clear but has perhaps to do with the truncation of the share price. Again we do not consider these to be legitimate for the purpose of approximation.




Figure 3: The upper panel shows the daily returns of the BMW share price. The lower panel shows the values near zero and a large atom at zero.

We make three additional comments. Firstly, although the idea of approximation

presented above is in terms of simulations these are not always necessary as the definition of “looks like” can sometimes be derived analytically. Secondly, if the parameter space is infinitely dimensional as in non-parametric regression the idea of approximation has to be augmented by considerations of simplicity. We refer to Davies and Kovac (2001) and (2004). Thirdly, in the above the degree of closeness of a probability model to the data is measured in a certain sense by the number α which specifies the word “typical”. In some cases what is statistically relevant may be irrelevant within the conceptual framework within which the data were gathered. A peak may be statistically significant but irrelevant for the physicist or engineer who produced the data. In such case the measure of closeness may be one which is not expressed in terms of probability.

3 Approximating Financial Data

3.1 Quantiles and long range memory

As argued above it will not be sufficient for a model to simply reproduce the stylized facts and more attention to detail is necessary. If we consider the stylized fact of heavy tails then we will not only require that the model reproduce the heavy tails but also that simulated quantiles are close to the quantiles of the data set we are seeking to model. Careful thought about how this may best be done is necessary. We may make a conscious decision not to model any zeros which the real data may exhibit (Figure 3). We may also make a conscious decision not to model the outliers which are present in the Standard and Poor’s data (Figure 1). This may be done by restricting attention to the 0 – 0.99–quantiles and we shall now do so. In the following we shall treat the positive and negative returns on the same basis and consider only the quantiles of the absolute values of the returns. Using the parameters (3) we can simulate samples of the same size $n = 9558$ as the Standard and Poor’s data and for each sample we can calculate a quantile curve $(p, Q(p)), p = 0, \dots, 0.99$. On the basis of 10000 simulations we calculate lower and upper bounds so that 98% of all simulated quantile curves lie between the two bounds. The upper panel of Figure 4 shows the results together with the corresponding quantile curve (denoted by stars) for the Standard and Poor’s data. Because of the scale of the upper panel it is not possible to distinguish between the lower bound and the Standard and Poor’s quantiles. For values of p up to about 0.7 the Standard and Poor’s curve lies below the lower bound as is shown in the lower panel of Figure 4 which plots the difference between the two curves. In this sense the quantiles of the GARCH process are somewhat too large. More relevant is the fact that many GARCH simulations results in very large quantiles and a model which does this may well be not acceptable. A similar situation holds for the autocorrelations as shown in Figure 5. The upper and lower bounds again contain about 98% of all simulated autocorrelation curves and the autocorrelation curve of the Standard and Poor’s process lies well between them. Again, a model which produces such large variations may not be regarded as satisfactory although it is by no means




Figure 4: The upper panel shows the lower and upper bounds for the quantile curves of a GARCH process with parameters given by (3). About 98% of the simulated curves lie between the two bounds. The quantile curve for the Standard and Poor's data is shown by the starred line. The lower panel shows the difference between the quantile curve of the Standard and Poor's data and the lower bound.

easy failing any relevant theory or empirical evidence to decide on what degree of variation is acceptable.

3.2 The volatility of the volatility

Figure 6 allows a direct comparison of the Standard and Poor's data shown in the top panel of Figure 1 and the GARCH simulation in the top panel of Figure 2. To aid comparison the Standard and Poor's data has been truncated to the same scale. There are similarities and differences. One of the differences is the presence of outliers in the Standard and Poor's data which we have chosen to ignore. It is sometimes difficult to quantify perceived differences. It is easy to recognize a friend's photograph amongst many photographs of strangers but not easy to write a computer programme to perform this task automatically. Sometimes theory may help to specify relevant differences but in its absence it is a matter of trial and error. A close visual inspection of the two processes




Figure 5: The panel shows the upper and lower bounds for autocorrelations of the absolute values of a GARCH process with parameters given by (3). About 98% of all autocorrelation functions lie between these two bounds. The starred curve gives the corresponding autocorrelations for the Standard and Poor’s data.

indicates, at least to the author, that the Standard and Poor’s data exhibits more frequent changes in volatility, the volatility of the volatility, than the simulated data. As the volatility of the volatility may be a relevant factor we now attempt to quantify it. We take a non-parametric approach. Applying the model (1) to real data (denoted by lower case letters) we look for a decomposition of the form

$$r(i) = \sigma(i)z(i) \quad (4)$$

where the $r(i)$ are the recorded daily returns. The $\sigma(i)$ and $z(i)$ are confounded but in accordance with the model (1) we intend to choose the $\sigma(i)$ so that the $z(i)$ “looks like” standard Gaussian noise. We make this precise on noting that (1) implies

$$\sum_{i \in I} \frac{R(i)^2}{\Sigma(i)^2} = \sum_{i \in I} Z(i)^2 \stackrel{D}{=} \chi^2(|I|) \quad (5)$$

where $\chi^2(k)$ denotes a chi-squared random variable with k degrees of freedom and $|I|$ denotes the number of points i in the interval I . As (5) holds for every interval I we are lead to the following system of inequalities

$$\text{qu}((1 - \alpha_n)/2, |I|) \leq \sum_{i \in I} \frac{r(i)^2}{\sigma(i)^2} \leq \text{qu}((1 + \alpha_n)/2, |I|), \quad I \in \mathcal{I}. \quad (6)$$

where $\text{qu}(\alpha, k)$ denotes the α -quantile of a chi-squared random variable with k degrees of freedom and \mathcal{I} is a family of intervals I . The default choice of α_n is

$$\alpha_n = 1 - 2 \exp(-1.15 \log(n)) / \sqrt{4.3\pi \log(n)} \quad (7)$$




Figure 6: The upper panel shows the simulated GARCH process of Figure 2. The lower panel shows the Standard and Poor's data of Figure 1 truncated to the size of the GARCH simulation.

which corresponds to the choice $\sqrt{2.3 \log(n)}$ for the threshold in Davies and Kovac (2001). The inequalities (6) alone do not determine the $\sigma(i)$ and, as in other non-parametric problems, some form of regularization is required. The one we choose is to take $\sigma(t)$ to be piecewise constant and then to minimize the number of intervals of constancy subject to (6). This leads to an optimization problem for the data at hand (Davies and Kovac 2001 and 2004) whose exact solution is not algorithmically simple. We propose the following procedure. We start with the first observation $r(1)$ and note that the volatility σ_1 for the first interval of constancy satisfies the inequalities

$$\sigma_l(1) \leq \sigma(1) \leq \sigma_u(1)$$

where the lower and bounds $\sigma_l(1)$ and $\sigma_u(1)$ are given by

$$\sigma_u(1) = \sqrt{r(1)^2 / \text{qu}((1 - \alpha_n)/2, 1)} \quad (8)$$

$$\sigma_l(1) = \sqrt{r(1)^2 / \text{qu}((1 + \alpha_n)/2, 1)}. \quad (9)$$

If the first interval contains t observations with lower and upper bounds $\sigma_l(t)$ and $\sigma_u(t)$ respectively the bounds for the first $t + 1$ observations are given by

$$\sigma_u(t+1)^2 = \min \left\{ \sigma_u(t)^2, \min_{1 \leq j \leq t+1} \left\{ \sum_{i=j}^{t+1} r(i)^2 / \text{qu}((1 - \alpha_n)/2, t + 2 - j) \right\} \right\} \quad (10)$$

$$\sigma_l(t+1)^2 = \max \left\{ \sigma_l(t)^2, \max_{1 \leq j \leq t+1} \left\{ \sum_{i=j}^{t+1} r(i)^2 / \text{qu}((1 + \alpha_n)/2, t + 2 - j) \right\} \right\}. \quad (11)$$

If for some t we have $\sigma_u(t) < \sigma_l(t)$ then the inequalities cannot be satisfied for a constant volatility and we start with a new interval. In practice we are somewhat more restrictive as, for reasons of interpretability, we wish the squared volatility $\sigma(I)^2$ over an interval I of constancy to be the mean of the squared returns over that interval, $\sigma(I)^2 = \sum_{i \in I} r(i)^2 / |I|$. Because of this we check at each stage the mean lies between the the squares of the lower and upper bounds

$$\sigma_l(t)^2 \leq \frac{1}{|I|} \sum_{i \in I} r(i)^2 \leq \sigma_u(t)^2. \quad (12)$$

and if this is not so we start with a new interval.

Using this procedure we can quantify the volatility of the volatility by the number of intervals of constancy and also by the distribution function of the lengths of the intervals, that is, the sojourn times at the different volatilities. In the case of Gaussian white noise for sample sizes from $n = 500$ to 10000 the procedure results in one interval of constancy in about 60% of the cases and the mean number of intervals is about 1.6.

The top panel of Figure 7 shows the Standard and Poor's data with the 49 intervals of constancy. The centre panel shows the observations 2001-3000. The bottom panel shows the distribution function of the sojourn times at a constant volatility. The shorter these are, the more often the volatility changes. The Standard and Poor's data results in 49 intervals of constancy. We can now use this quantification of the volatility of the volatility can be reproduced by the GARCH model. The 98% bounds for the numbers of intervals of constancy are 27–43 with a mean value of 34.95. This confirms the impression that the volatility of the volatility for the GARCH model is too low. Figure 8 shows the 98% bounds for the distribution function of the sojourn times together with the distribution function for the Standard and Poor's data. It is seen that it does not lie within the bounds so again the model is not able to reproduce this aspect of the data.

Mercurio and Spokoiny (2000) also give a non-parametric piecewise constant volatility analysis of financial data based on the idea of adaptive weighted




Figure 7: The top panel shows the Standard and Poor's data with 49 intervals of constant volatility. The centre panel shows the observations 2001–3000. The bottom panels shows the distribution function of the sojourn times.

8h]




Figure 8: The upper and lower bounds for the distribution function of the sojourn times for the GARCH process and the distribution function of the sojourn times for the Standard and Poor's data (starred line).

smoothing (see Polzehl and Spokoiny 2002). Although their method is conceptually different and based on local behaviour rather than some form of global regularization the results are very similar. Peters (2003) uses a piecewise constant volatility model based on a piecewise linearization of the quadratic variation process. He applies the method to high frequency data whereas we use only daily data.

3.3 Combining the features

The features defined by the quantiles of the absolute returns, the autocorrelation function of the absolute returns, the number of intervals of constant volatility and the distribution of the sojourn times at the volatility levels may be combined in that we require a typical sample generated under the model to reproduce all four features. We have done this using a probabilistic definition of approximation and spent $0.02=1-0.98$ of probability on each feature. This leads to a value of the *alpha* which quantifies “typical” of at least $1 - 0.08 = 0.92$. The features are however related and the actual value of α will probably exceed the lower bound of 0.92. On the basis of 10000 simulations a more accurate value for α for the Standard and Poor's data is 0.947. A model will be an adequate approximation if the corresponding quantities for the real data lie within the specified bounds. The results for the Standard and Poor's data and the GARCH model with parameters given by (3) is that only the autocorrelation function of the Standard and Poor's data lies within the corresponding bounds. We conclude that the model is not an adequate approximation. The results for the German motor-car company BMW are similar. The maximum likelihood parameter values for the

GARCH model are

$$\alpha_0 = 0.0579, \alpha_1 = 0.8473, \beta_1 = 0.1327 \quad (13)$$

and it turns out that none of the four features lies within the bounds. The limits for the number of intervals for example are 45 and 66 whereas the BMW data have 73 intervals.

Other models for such time series include COGARCH, continuous GARCH, non-Gaussian Ornstein-Uhlenbeck models, cascade models, non-linear adaptive systems models and genetic adaptive learning models have also been proposed. We refer to Embrechts, Klüppelberg and Mikosch 1997, Mikosch and Stărciă 1999, Barndorff-Nielsen and Shephard (2001) and (2002), Hommes (2002), Lux and Schornstein (2002), Mikosch 2003, Lux 2003, Klüppelberg, Lindner and Maller (2004) and (2005) and the references given there. Initial indications are that these models cannot reproduce all the four features simultaneously but an exhaustive investigation remains to be carried out. In the next section we propose another model which does seem able to reproduce these features at least for some financial data series.

4 A Stochastic Model for Volatility

4.1 A finer scale for the volatility

In the model (1) the volatility $\Sigma(t)$ and the white noise process $Z(t)$ are confounded. They can be separated by proposing a model for $\Sigma(t)$ such as a GARCH model or, as in the last section, by regularizing $\Sigma(t)$ in some manner. Specific models differ by their specification of the volatility process $\Sigma(t)$ and we now use non-parametric piecewise constant volatility idea to help in the construction of a model for the volatility. If we replace the α_n of (7) by $\alpha = 0.99$ then for the Standard and Poor's data we obtain 453 intervals of constant volatility. Figure 9 shows the plot of their logarithms. It suggests that to a first approximation the log-volatility can be modelled by a stationary process with long range memory. Figure 10 plots the sojourn times against the volatility level for the volatility process of Figure 9. It is clear that the sojourn times are small both for high and low volatilities.

4.2 A stochastic model for the volatility

We start with a stationary process X to model the log-volatility. We take X to have a spectral density of the form

$$f(\omega) = (1 - \alpha)\lambda_1 \exp(-\lambda_1\omega) + \alpha\lambda_2 \exp(-\lambda_2\omega) + \delta, \quad 0 \leq \omega \leq 2\pi, \quad (14)$$

where $\lambda_1, \lambda_2, \delta \geq 0$ and $\alpha, 0 \leq \alpha \leq 1$ are parameters. The initial volatility process is

$$\Sigma_0(t) = \exp(\sigma X(t)) \quad (15)$$




Figure 9: The upper panel shows the logarithm of the piecewise constant volatilities for the Standard and Poor's data with $\alpha_n = 0.99$. The lower panel shows the observations 1:1000.

where $\sigma > 0$ is a further parameter. We model the sojourn times at a particular volatility level as follows. Given a sample $\Sigma_0(i), i = 1, \dots, m$ we put

$$\tau(i) = \lfloor c\Sigma_0(i)E(i)/(1 + \exp(b\Sigma_0(i))) \rfloor \quad (16)$$

where b and c are further parameters and the $E(i)$ are independently and exponentially distributed random variables with mean 1. The functional form of (16) is suggested by Figure 10 where the sojourn time is small for small and large volatilities. The volatility is initially $\Sigma_0(1)$ and remains at this level for a time $\tau(1)$ after which it changes to $\Sigma_0(2)$ and remains at this level for a time $\tau(2)$ and so forth. We denote this volatility by $\Sigma(t)$. The final process $R(t)$ is then given by (1) with $Z(t)$ standard Gaussian white noise. This model has in all seven parameters, $\lambda_1, \lambda_2, \delta, \alpha, \sigma, b$ and c .

4.3 Model choice

Given data $r(1), \dots, r(n)$ and the model of the previous section we must now determine which parameter values if any give an adequate model for the data. As argued above this means that simulated data should “look like” the real data in certain respects which are judged to be relevant. The ones we have chosen are the marginal distributions of the returns, the autocorrelations of the absolute




Figure 10: Standard and Poor’s sojourn times plotted against volatility with $\alpha = 0.99$ upto a standardized (median 1) volatility of 10.

terms, the number of intervals of constant volatility and the distribution of the sojourn times Section 3.2. The parameter values we use for the Standard and Poor’s data are

$$\lambda_1 = 10, \lambda_2 = 600, \alpha = 0.60, \delta = 0.4, \sigma = 0.47, b = 0.5, c = 100. \quad (17)$$

Figure 11 shows the results for the quantiles of the absolute returns, the autocorrelations and the distribution function of the sojourn times of constant volatility. As for the GARCH model the bounds are chosen so that under the model the values for the simulated data lie within the bounds with probability about 0.98 for each feature separately. The results for the number of intervals give an upper bound of 62, a lower bound of 38 and a mean value of 50.40 compared with 49 intervals for the Standard and Poor’s data. It is seen that the model of Section 4.2 with the parameter value (17) can reproduce these features of the Standard and Poor’s data and in this precise sense it is an adequate approximation. The value of α which quantifies the word “typical” is at least 0.92. Simulations give a more accurate value of 0.945.

The number of intervals of constant volatility for the BMW data is 73 which is well outside the limits 49–62 for the Standard and Poor’s data. This shows that the parameter values (17) do not give an adequate approximation for the BMW data. The set of parameter values $\mathcal{A}(\mathbf{x}_n)$ which are an adequate approximation for a data set is called an approximation region in Davies (1995). Although it has certain similarities with, and occasionally coincides with, a confidence interval it is somewhat different. An adequacy region can be empty which means that there




Figure 11: Model of Section 4.2: The top panel shows the bounds for the quantiles of the absolute returns, the mean values and the values for the Standard and Poor's data. The centre panel shows the the bounds for the autocorrelations of the absolute returns, the mean autocorrelations and the autocorrelations for the Standard and Poor's data. The bottom panel shows the distribution function of the sojourn times, the mean distribution function and the distribution function for the Standard and Poor's data.

is no adequate model in the models class under consideration. Any reasonable procedure fro model choice must include this option.

References

- [1] O. E. Barndorff-Nielsen and N. Shephard. Non-Gaussian Ornstein-Uhlenbeck models and some of their uses in financial economics. *Journal of Royal Statistical Society Series B*, 63:167–241, 2001.
- [2] O. E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of Royal Statistical Society Series B*, 64:253–280, 2002.

- [3] M. Cohen. The background of configural polysampling: a historical perspective. In S. Morgenthaler and J. W. Tukey, editors, *Configural Polysampling: A Route to Practical Robustness*, chapter 2. Wiley, New York, 1991.
- [4] P. L. Davies. Data features. *Statistica Neerlandica*, 49:185–245, 1995.
- [5] P. L. Davies and A. Kovac. Local extremes, runs, strings and multiresolution (with discussion). *Annals of Statistics*, 29(1):1–65, 2001.
- [6] P. L. Davies and A. Kovac. Densities, spectral densities and modality. *Annals of Statistics*, 32(3):1093–1136, 2004.
- [7] P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin, 1997.
- [8] C. H. Hommes. Modeling the stylized facts in finance through simple non-linear adaptive systems. *Proceedings of the National Academy of Sciences*, 99:7221–7228, 2002.
- [9] N. Kaldor. Capital accumulation and economic growth. In F.A. Lutz and D.C. Hague, editors, *The Theory of Capital*. St. Martin’s Press, New York, 1961.
- [10] C. Klüppelberg, A. Lindner, and R. Maller. A continuous-time GARCH process driven by a Lévy process: Stationarity and second-order behaviour. *Journal of Applied Probability*, 41, 2004.
- [11] C. Klüppelberg, A. Lindner, and R. Maller. Continuous time volatility modelling: COGARCH versus Ornstein-Uhlenbeck models. In Y. Kabanov, R. Lipster, and J. Stoyanov, editors, *From Stochastic Calculus to Mathematical Finance, The Shiryaev Festschrift*. Springer, to appear.
- [12] T. Lux. The multi-fractal model of asset returns: its estimation via gmm and its use for volatility forecasting. Technical report, Department of Economics, University of Kiel, Germany, 2003.
- [13] T. Lux and S. Schornstein. Genetic learning as an explanation of stylized facts of foreign exchange markets. Technical Report 29/02, Economic Research Centre of the Deutsche Bundesbank, 2002.
- [14] D. Mercurio and V. Spokoiny. Statistical inference for time-inhomogeneous volatility models. Technical Report WIAS-Preprint 583, Weierstrass Institute for Applied Mathematics and Stochastics, Mohrenstraße 39, 10117 Berlin, Germany, 2000.
- [15] T. Mikosch and C. Stărică. Change of structure in financial time series, long range dependence and the garch model. Technical report, Laboratory of Actuarial Mathematics, University of Copenhagen, Denmark, 1999.

- [16] T. Mikosh. Modeling dependence and tails of financial time series. In B. Finkenstädt and H. Rootzen, editors, *Extreme Values in Finance, Telecommunications and the Environment*, number 99 in Monographs on Statistics and Applied Probability, pages 185–286. Chapman and Hall, 2003.
- [17] R. T. Peters. Shocks in the S&P500. Technical report, Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Plantage Muidergracht 24 1018TV Amsterdam, The Netherlands, 2003.
- [18] J. Polzehl and V. Spokoiny. Local likelihood modeling by adaptive weights smoothing. Technical Report WIAS-Preprint 787, Weierstrass Institute for Applied Mathematics and Stochastics, Mohrenstraße 39, 10117 Berlin, Germany, 2002.
- [19] J. W. Tukey. Issues relevant to an honest account of data-based inference, partially in the light of Laurie Davies’s paper. Princeton University, Princeton, 1993.