On testing the adequacy of stable processes under conditional heteroscedasticity

Rohit S. Deo [*]

Abstract

We consider a recently proposed method of estimating the tail index and testing the goodness-of-fit of dependent stable processes. Through Monte Carlo simulations, we evaluate the ability of the procedure to distinguish between stable and non-stable processes in the presence of non-linear dependence and to estimate the tail index of the distribution. We then apply the test to black market East European exchange rates, whose distributional and tail behaviour has been analysed previously in the literature. After adjusting for seasonality, we conclude, unlike the earlier analysis, that a stable process cannot be rejected as a model for some of the currencies. Estimates of the tail index for these currencies are also obtained. © 2002 Elsevier Science B.V. All rights reserved.

1. Introduction

It is now well accepted in the literature that the empirical distributions of exchange rates exhibit thicker tails than are to be expected from Gaussian distributions. However, there does not seem to be any agreement as to which distribution is the appropriate one to model this behaviour. Mandelbrot (1963) and

[*] Tel.: +1-212-998-0469; fax: +1-212-995-4003.
E-mail address: rdeo@stern.nyu.edu (R.S. Deo).

Fama (1963), in their seminal papers, proposed the non-Gaussian stable distributions as candidate distributions for such data. Since then, other researchers have proposed alternatives like the Student's t distribution and mixtures of normal distributions. Unfortunately, so far there is no simple goodness-of-fit test in the literature for any of these hypotheses and so disagreement about the underlying distribution remains. To the best of our knowledge, none of the studies so far on the applicability of stable distributions to financial data have carried out formal goodness-of-fit tests. As Baillie (1993) points out, "One apparently appealing property of stable distributions is the invariance of the characteristic exponent over changes in the sampling interval. As mentioned by Mittnik and Rachev, this property is generally not consistent with previous empirical findings. However, the estimation of the characteristic exponent has invariably not been by maximum likelihood methods and the standard errors on the estimated exponent parameter may well be very wide. No formal test of this property of stable distributions appears to have yet been done." A further complication which has been generally ignored in most empirical studies of exchange rate distributions which deal with stable distributions is the lack of independence in such data. It is a well-accepted fact that though the returns themselves are uncorrelated, their absolute values or squares are not and hence the returns are not independently distributed. Any procedure which attempts to check the adequacy of a marginal probability distribution to the data must at least make some attempt to incorporate this non-linear dependence.

Recently, Deo (2000) has proposed a goodness-of-fit test for stable distributions which remains valid when the data is m-dependent, where m is a finite unknown integer. (A series $\{X_t\}$ is said to be m-dependent if X_i and X_j are independent for all i,j such that $|i-j| > m$). Thus, the test allows for possible non-linear dependence in the series. Though m-dependence might be a simplification of the true underlying dependence, we believe that it is far more realistic than the assumption of independence made by all studies involving stable distributions as models of marginal distributions so far. Furthermore, the test does not assume any knowledge about the value of m beyond that it is finite. Thus, m may be large enough to provide an adequate approximation to the true underlying process.

In this paper we examine the performance of this goodness-of-fit test to distinguish between stable and non-stable processes which are conditionally heteroscedastic and also dependent at all lags. We then apply the testing procedure to black market East European exchange rates, which have been analysed in the literature. We find that the exchange rates have seasonality. After adjusting for seasonality, our goodness-of-fit test shows that a marginal stable distribution cannot be rejected as a model for some of the returns.

The layout of this paper is as follows. In Section 2, we describe the estimation and goodness-of-fit testing procedure of Deo (2000) and report the Monte Carlo simulation results. In Section 3, the empirical data analysis of the exchange rates is carried out.

2. The goodness-of-fit test

We briefly describe the intuition behind the estimation and goodness-of-fit procedure of Deo (2000) for stable distributions. Let X_1, X_2, ..., X_n be n independently and identically distributed observations from a stable distribution with tail index α. Using the form of the characteristic function of stable distributions (see Samorodnitsky and Taqqu, 1994), it can be shown that for integers k and s such that $k > s$,

$$X_1 + X_2 + \ldots + X_k - X_{k+1} - X_{k+2} - \ldots - X_{2k} \overset{D}{\underset{\sim}{=}} \left(\frac{k}{s}\right)^{1/\alpha}$$

$$\times (X_1 + \ldots + X_s - X_{s+1} - X_{s+2} - \ldots - X_{2s}), \tag{1}$$

where $\underset{\sim}{D}$ denotes "identically distributed". This relationship arises basically due to the fundamental property of invariance of stable distributions under aggregation. From Eq. (1), it follows that

$$\left[\left(\log\left(\frac{k}{s}\right)\right)^{-1} \log \frac{|X_1 + X_2 + \ldots + X_k - X_{k+1} - X_{k+2} - \ldots - X_{2k}|}{|X_1 + \ldots + X_s - X_{s+1} - X_{s+2} - \ldots - X_{2s}|}\right], \tag{2}$$

will be unbiased for α^{-1}. There is nothing special about the set $(X_1, X_2, \ldots, X_{2k})$ in Eq. (2) and one may use any arbitrary group $(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$ instead. Thus, averaging the form (2) over all subsets $(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$ would give an estimator of α^{-1}, and hence of α. This is the rationale behind the estimation procedure for α proposed in Deo (2000). Furthermore, by varying k and s in Eq. (2), one gets several different estimators of the same quantity α. Once again, these different estimators of the tail index α arise from the fundamental property of invariance of stable distributions under aggregation. From a practical point of view, computing an estimator of α by averaging Eq. (2) over all possible subsets $(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$ is a computationally prohibitive task. As a result, Deo (2000) proposed using incomplete estimators of α, which we describe next.

For a positive integer $k > 1$, let $A = \{(i_1, i_2, \ldots, i_{2k}): 1 \le i_1 < i_2 < \ldots < i_{2k} \le n\}$ be the collection of the $\binom{n}{2k}$ $2k$-tuples obtained from the set $\{1, 2, \ldots, n\}$. For a positive integer N_0, let D be the collection of N_0 elements of A obtained by sampling with replacement from A. Let s and k be positive integers such that $1 < 2s < 2k < n$ and define

$$\phi_{k,s}^*(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$$

$$= \left\{\log\left(\frac{k}{s}\right)\right\}^{-1} \left[\log \frac{|X_{i_1} + \ldots + X_{i_k} - X_{i_{k+1}} - \ldots - X_{i_{2k}}|}{|X_{i_1} + \ldots + X_{i_s} - X_{i_{s+1}} - \ldots - X_{i_{2s}}|}\right]. \tag{3}$$

Now let

$$\hat{\alpha}_{k,s} = \left[N_0^{-1} \sum_D \phi_{k,s}\left(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}} \right) \right]^{-1}, \tag{4}$$

where $\phi_{k,s}(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$ is the symmetric version of the function $\phi_{k,s}^*(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$, obtained by

$$\phi_{k,s}\left(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}} \right) = \frac{1}{(2k)!} \sum_p \phi_{k,s}^*\left(X_{j_1}, X_{j_2}, \ldots, X_{j_{2k}} \right),$$

where $\sum_p$ denotes summation over all the $(2k)!$ permutations $(j_1, j_2, \ldots, j_{2k})$ of $(i_1, i_2, \ldots, i_{2k})$. The estimator (4) is *incomplete* in that it no longer averages over all possible choices $(X_{i_1}, X_{i_2}, \ldots, X_{i_{2k}})$ but over a smaller selection N_0 of such choices. Note that for a fixed value of k, the set D remains the same for all values of $s < k$ and hence the set D does not have to be selected again for different values of s. Deo (2000) showed that under the hypothesis that the observations $(X_1, \ldots, X_n)$ come from an m-dependent stable process, as $N_0 \to \infty$,

$$\sqrt{n}\left(\hat{\boldsymbol{\alpha}} - \alpha \mathbf{1} \right) \overset{D}{\to} N(\mathbf{0}, \boldsymbol{\Sigma}),$$

where $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}_{k,s_1}, \ldots, \hat{\alpha}_{k,s_p})'$, $\mathbf{1} = (1, \ldots, 1)'$, $s_1 < \ldots < s_p < k$ and $\boldsymbol{\Sigma}$ is as in Theorem 1 of Deo (2000). Furthermore, if $N_0^{-1} = O(n^{-2})$, a consistent estimator of $\boldsymbol{\Sigma}$, the asymptotic variance of $\hat{\alpha}_{k,s}$, can be obtained, as provided in Eq. (19) of Deo (2000).

Note that neither the construction of the estimator nor the limiting results require knowledge of m, which governs the degree of dependence, nor is any assumption made about the scale, location and skewness parameters of the distribution of $\{X_t\}$. As a matter of fact, it is easy to see that the estimators are scale and location invariant. Furthermore, by changing values of k and s, various estimators of the same quantity α may be obtained. All these estimators are jointly asymptotically normal and their limiting variance covariance matrix can be estimated. Hence, a goodness-of-fit test of stability may be obtained by computing a t-statistic based on the difference of two different estimates of α. For example, one could compute

$$t = \frac{\hat{\alpha}_{k,s_i} - \hat{\alpha}_{k,s_j}}{\text{s.e.}\left(\hat{\alpha}_{k,s_i} - \hat{\alpha}_{k,s_j} \right)} \quad 1 < i < j < k, \tag{5}$$

where s.e.$(\hat{\alpha}_{k,s_i} - \hat{\alpha}_{k,s_j})$ is the standard error of $\hat{\alpha}_{k,s_i} - \hat{\alpha}_{k,s_j}$ and is obtained from the estimate of the limiting variance matrix $\boldsymbol{\Sigma}$. Under the null hypothesis that the series $\{X_t\}$ comes from an m-dependent process with a marginal stable distribution, the t-statistic in Eq. (5) will have an asymptotic standard normal distribution.

Though the procedure outlined above assumes m-dependence in the series, this assumption may be unrealistic in real data. To assess the performance of the

procedure under infinite dependence, we carried out a simulation study. To study the size properties, we generated 1000 observations of the series

$$X_t = Y_t^{1/2} e_t, \tag{6}$$

where $Y_t = \tau Y_{t-1} + v_t$. The series $\{e_t\}$ was chosen to be *i.i.d.* $N(0,1)$ and independent of the innovation series $\{v_t\}$. The series $\{v_t\}$ was a sequence of independent positive stable random variables chosen in such a way that the marginal distribution of $\{X_t\}$ was stable with tail index α. The process (6) was originally proposed in de Vries (1991) as an example of a conditionally heteroscedastic process with a marginal stable distribution. We generated the series $\{X_t\}$ for two values of α, 1.7 and 1.8. These values were chosen since estimated values of the tail index in most empirical studies of exchange rates range between 1.6 and 1.85. For each fixed value of α, we generated observations on $\{X_t\}$ for values of $\tau = 0.2, 0.3, 0.4, 0.5$ and 0.7. We selected these values of τ since de Vries (1991) had estimated values of τ between 0.15 and 0.46 for the same model (Eq. (6)) using European exchange rate returns. For each pair of values (α, τ), we computed three estimators $\hat{\alpha}_{3,2}$, $\hat{\alpha}_{3,1}$ and $\hat{\alpha}_{2,1}$ defined by Eq. (4) above. The value of N_0 was chosen to be 4,000,000 based on simulations in Deo (2000). Each experiment was replicated 1000 times.

In Table 1, we present the averages of each of the three estimators of α over the 1000 replications for each pair of values (α, τ). It is seen that for a fixed estimator, the bias increases as the value of τ increases. This is to be expected, since the degree of dependence in the series increases with τ. However, it is reassuring that even in the worst possible case, when $\alpha = 1.7$ and $\tau = 0.7$, the bias is never more than 1%. Indeed, when $\tau < 0.7$, the bias in all three estimators of α is less than 0.5% for both values of α. Furthermore, for every pair of values of (α, τ), the bias in all three estimators is the least in $\hat{\alpha}_{2,1}$. This superiority of $\hat{\alpha}_{2,1}$ to the other two estimators when $\alpha \leq 1.8$ was also noted in the Monte Carlo study in Deo (2000) in the case of independent observations. As a matter of fact, it was shown in that study that $\hat{\alpha}_{2,1}$ uniformly outperforms, by a wide margin, the regression estimator which is currently the most popular estimator of the tail index. See Akgiray and Lamoureux (1989).

Table 1
Averages of estimators over 1000 replications for Stable Process Standard deviations are in parentheses

τ	$\alpha = 1.7$			$\alpha = 1.8$		
	$\hat{\alpha}_{3,2}$	$\hat{\alpha}_{3,1}$	$\hat{\alpha}_{2,1}$	$\hat{\alpha}_{3,2}$	$\hat{\alpha}_{3,1}$	$\hat{\alpha}_{2,1}$
0.2	1.704 (0.079)	1.703 (0.069)	1.703 (0.065)	1.803 (0.068)	1.802 (0.061)	1.801 (0.058)
0.3	1.705 (0.086)	1.704 (0.076)	1.704 (0.071)	1.804 (0.074)	1.803 (0.066)	1.802 (0.063)
0.5	1.709 (0.104)	1.707 (0.092)	1.706 (0.086)	1.807 (0.090)	1.805 (0.080)	1.804 (0.075)
0.7	1.717 (0.130)	1.713 (0.115)	1.711 (0.107)	1.812 (0.111)	1.809 (0.099)	1.807 (0.093)

Table 2
Monte Carlo sizes for goodness-of-fit tests at the 5% l.o.s.

τ	$\alpha = 1.7$			$\alpha = 1.8$		
	Z_1	Z_2	Z_3	Z_1	Z_2	Z_3
0.2	0.049	0.046	0.016	0.045	0.034	0.009
0.3	0.062	0.048	0.018	0.041	0.033	0.031
0.5	0.098	0.088	0.039	0.057	0.051	0.017
0.7	0.160	0.149	0.068	0.100	0.085	0.032

We also assessed the performance of three goodness-of-fit tests in our study. The test statistics are based on t-statistics of the differences between two estimators of α, as defined in Eq. (5). These tests are: (i) the test based on the difference $\hat{\alpha}_{3,2} - \hat{\alpha}_{3,1}$ and denoted by Z_1; (ii) the test based on the difference $\hat{\alpha}_{3,2} - \hat{\alpha}_{2,1}$ and denoted by Z_2; and (iii) the test based on the difference $\hat{\alpha}_{3,1} - \hat{\alpha}_{2,1}$ and denoted by Z_3. The test statistics all have limiting standard normal distributions. The variance covariance matrix of the vector $\hat{\boldsymbol{\alpha}}(\hat{\alpha}_{3,2}, \hat{\alpha}_{3,1}, \hat{\alpha}_{2,1})'$ was calculated by using the prewhitened kernel estimator described in Eq. (19) of Deo (2000). In Table 2, we report the empirical sizes of these three goodness-of-fit test statistics at the 5% level of significance for each pair of values of (α, τ). It is seen that for both values of α, the sizes of the tests increase as τ increases with the tests based on Z_1 and Z_2 becoming quite oversized at $\tau = 0.7$. The test based on Z_3 is most undersized for all values of τ, the loss in size increasing as τ decreases. For values of τ between 0.2 and 0.5, the test based on Z_3 seems to maintain its size reasonably well.

We next evaluated the performance of the procedure when the data was conditionally heteroscedastic but not stably distributed. Towards this end, we simulated 1000 observations from a GARCH(1,1) and an IGARCH(1,1) process. The GARCH(1,1) was generated as

$$X_t = \sigma_t e_t, \tag{7}$$

where $\sigma_t^2 = 4.67 \times 10^{-7} + 0.0529 X_{t-1}^2 + 0.915 \sigma_{t-1}^2$ and the e_t were $i.i.d.$ from a t-distribution with 6 df. The values of the parameters in the volatility process

Table 3
Averages of estimators over 1000 replications for GARCH(1,1) and IGARCH(1,1) Standard deviations are in parentheses

	GARCH(1,1)	IGARCH(1,1)
$\hat{\alpha}_{3,2}$	1.773 (0.116)	1.906 (0.081)
$\hat{\alpha}_{3,1}$	1.717 (0.124)	1.867 (0.107)
$\hat{\alpha}_{2,1}$	1.687 (0.128)	1.845 (0.122)

Table 4
Power at the 5% l.o.s. for GARCH(1,1) and IGARCH(1,1)

	GARCH(1,1)	IGARCH(1,1)
$\hat{\alpha}_{3,2}$	0.981	0.638
$\hat{\alpha}_{3,1}$	0.976	0.600
$\hat{\alpha}_{2,1}$	0.871	0.440

σ_t^2 were chosen to reflect those obtained when such models are fit to real data. The use of a *t*-distribution with 6 *df* for the innovations e_t implies that X_t has a thick-tailed distribution with tail index between 2 and 3. Hence, the process has a finite variance. In Table 3, we report the averages of each of the three estimators of α over 1000 replications for the GARCH(1,1) process. It is seen that the estimators are all biased downwards and actually indicate that the process has infinite variance. However, it should be noted that the estimators $\hat{\alpha}_{k,s}$ described above are designed for stable distributions and not for any arbitrary thick tailed distributions. Furthermore, it is seen from Table 4 that all the three goodness-of-fit tests soundly reject the process as being stably distributed. Hence, the estimation and testing procedure is able to detect the lack of a stable distribution for the data.

The IGARCH(1,1) process was also generated by Eq. (7), though now the volatility was modelled as $\sigma_t^2 = 4 \times 10^{-7} + 0.05 X_{t-1}^2 + 0.95 \sigma_{t-1}^2$ and the e_t were chosen to be *i.i.d.* $N(0,1)$. It is well known that such an IGARCH(1,1) process has infinite variance and it can be shown using results in Nelson (1990) that the tail index is approximately 1.99. In Table 3, we report the averages of each of the three estimators of α over 1000 replications. It is seen that the estimators reflect the fact that the process has infinite variance since the estimated tail index is less than 2. However, from Table 4 we see that once again the goodness-of-fit tests are able to detect quite well that the process is not stable. The power of the tests is lower than in the case of the GARCH(1,1) model. This is not too surprising and can be attributed to the fact that the IGARCH(1,1) process has infinite variance and is able to masquerade as a stable process more successfully than the finite variance GARCH(1,1).

The simulation study is encouraging and indicates that the estimation procedure is able to provide quite good estimates of the tail index when the process is stably distributed, even in the presence of strong conditional heteroscedasticity. Furthermore, the goodness-of-fit tests are capable of detecting non-stable processes, even when the process has infinite variance.

In the next section, we apply these methods to real data.

3. East European black market exchange rate data

The data we study are monthly black market exchange rate observations on seven East European currencies vis-a-vis the U.S. dollar for the period from

January 1955 to December 1990. The seven currencies are the Bulgarian Lev, the Czechoslovak Koruna, the East German Mark, the Hungarian Forint, the Polish Zloty, the Rumanian Lei and the Soviet Ruble. The data were studied in Koedijk and Kool (1992) and we have obtained them from their Appendix. The exchange rates for Bulgaria and the U.S.S.R. were modified since currency reforms took place in these countries in December 1961 and January 1961, respectively. In both countries, one new currency unit was set equal to 10 old currency units. We thus modified the exchange rates prior to the currency reforms to new currency units and used these in the subsequent analysis. This modification is the same one used in Koedijk and Kool (1992). The exchange rate returns, which are the variables we actually study, are defined as 100 times the first difference of the natural logarithms of the black market exchange rates.

Fig. 1 is a plot of the autocorrelation function of the Hungarian Forint returns. The plot reveals a roughly sinusoidal pattern, an indication of seasonality. This is confirmed by Fig. 2, which plots the periodogram of the Hungarian returns. The plot has a major spike corresponding to seasonality of period 12. The nature of the seasonal effect is illustrated in Fig. 3, which is a plot of the monthly 5% trimmed means (see Hoaglin et al., 1983, p. 311) against the corresponding months. The 5% trimmed means were used instead of the monthly means to reduce the effect of outliers in the data. The monthly average returns are lowest in the winter months
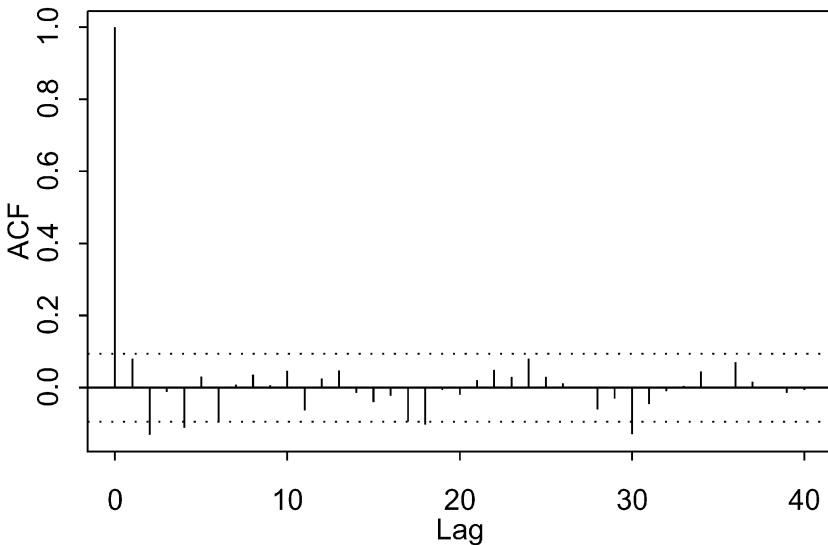
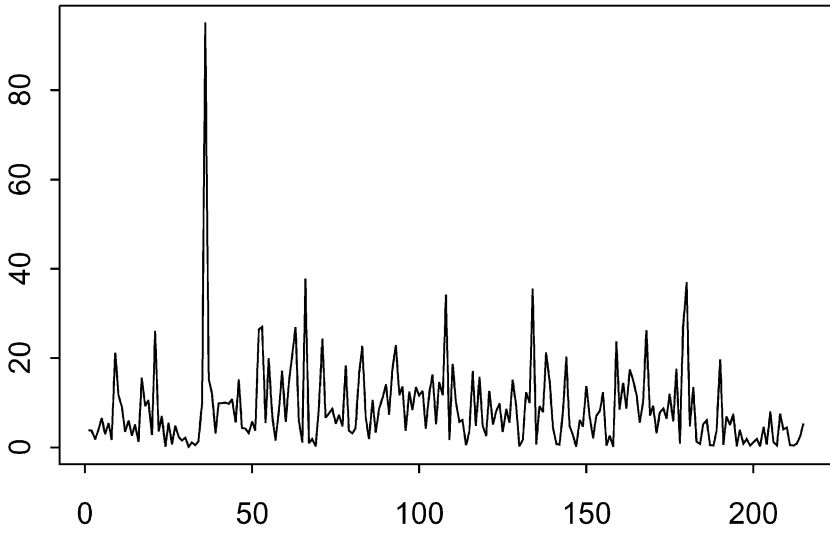# Series : forint

Fig. 1.

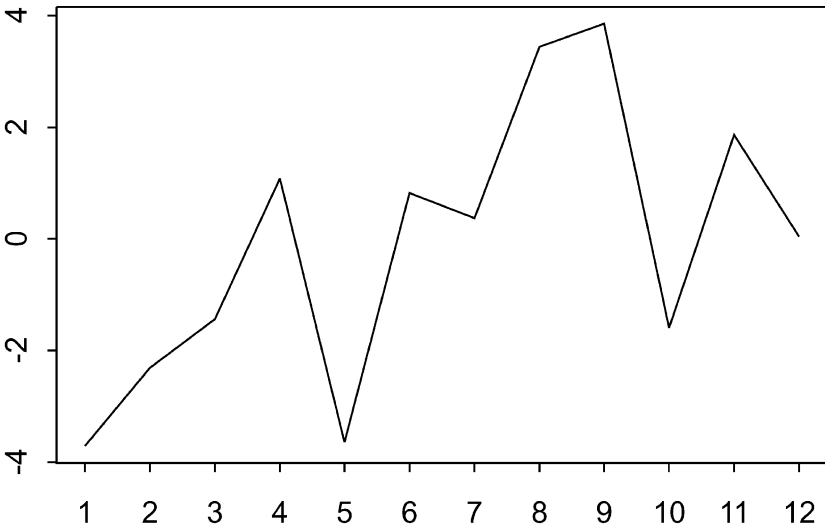# Periodogram: Raw Data

Fig. 2.

Trimmed Monthly Means

Fig. 3.

Table 5
p-Values for Fisher's test

	Lev	Koruna	Mark	Forint	Zloty	Lei	Ruble
Original data	0.000	0.000	0.052	0.004	0.179	0.000	0.899
Adjusted data	0.407	0.693	0.457	0.749	–	0.254	–

of January and February and peak in the summer months of August and September, probably an effect of the presence of tourists. There is a puzzling drop in the average in May, for which we do not have an explanation. Similar patterns reveal themselves in the autocorrelation plots, periodograms and monthly average plots of almost all the other currencies with monthly averages peaking in summer and dropping in winter. To check whether the spikes in the periodogram were significant, we carried out Fisher's test for periodicities (see Bloomfield, 1976, p. 112). The corresponding *p*-values are reported in Table 5 and indicate that there is overwhelming evidence of seasonality in all the currencies except the Polish Zloty and the Soviet Ruble. As a result, we de-seasonalised only the other five currencies by subtracting the monthly 5% trimmed averages from the corresponding observations (i.e., by subtracting the January trimmed average from the January observations, the February trimmed average from the February observations and so on). Fig. 4 is a plot of the periodogram of the de-seasonalised
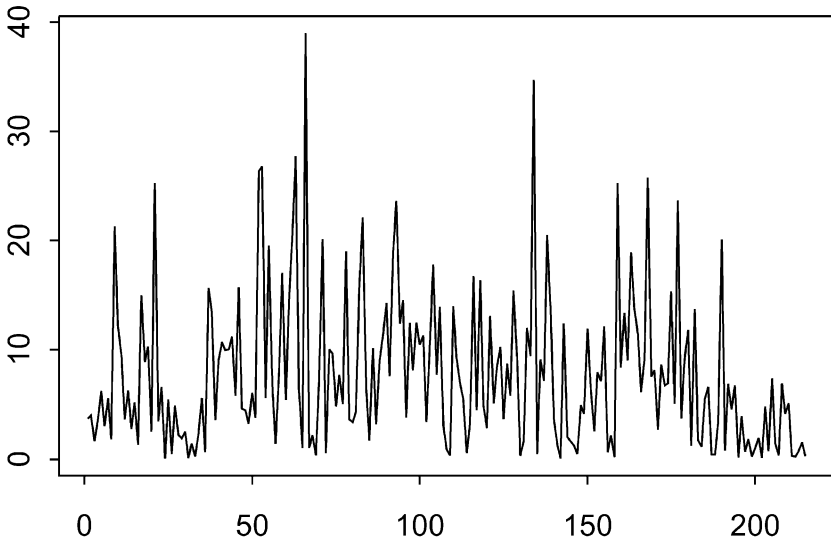

Fig. 4.

Hungarian returns and does not show a single significant spike. Similar behaviour is shown by the periodograms of the other de-seasonalised returns. The *p*-values of Fisher's test for periodicities for the de-seasonalised data are also shown in Table 5 and indicate that the seasonal component has indeed been successfully removed. The subsequent analysis was carried out on the de-seasonalised returns (except for the Polish Zloty and the Soviet Ruble), which we merely refer to as the returns henceforth.

For the sake of comparison, we fit using maximum likelihood estimation the benchmark GARCH(1,1) model

$$y_t = \sigma_t v_t$$
$$\sigma_t^2 = \alpha_0 + \alpha_1 y_{t-1}^2 + \beta_1 \sigma_{t-1}^2,$$

to the returns. Table 6 contains the estimated values of α_1 and β_1 when the innovations v_t were assumed to be standard normal variables. The columns headed *K* and *S* contain the coefficients of kurtosis and skewness for the standardized residuals from the fitted model, while the column headed BL_{15} contains the values of the Box–Ljung statistic for the squared standardized residuals at lag 15. Under the assumption of normality, the coefficient of kurtosis should be approximately normally distributed with mean 3 and variance $24/n$, while *S* should be approximately normally distributed with mean zero and variance $6/n$. We see that the GARCH(1,1) process with conditional normal distribution is clearly unable to account for the extreme kurtosis in the data. We then fit the GARCH(1,1) model with a conditional *t*-distribution to the data, where the degrees of freedom, v, of the *t*-distribution were estimated. Table 7 contains the estimated parameters of this model, with the standard error of the estimated degrees of freedom, $\hat{v}$, in parentheses. To facilitate comparison, the standardized residuals from the fitted model were transformed to standard normal variables by using first the inverse of the Student's *t* distribution and then the standard normal distribution function. The coefficients of kurtosis and skewness and the Box–Ljung statistic were then computed from these transformed residuals. It can be seen from Table 7 that the conditional *t*-distribution has successfully accounted for the kurtosis in the data.

Table 6
GARCH(1,1) with conditional normal distribution

	$\hat{\alpha}_1$	$\hat{\beta}_1$	K	S	BL_{15}
Bulgarian Lev	0.0265	0.9481	8.988	1.188	5.20
Czech Koruna	0.0829	0.8225	5.901	0.638	13.18
E. German Mark	0.1128	0.7666	16.119	1.603	5.23
Hungarian Forint	0.4931	0.0368	3.880	0.011	31.91
Polish Zloty	0.2296	0.7052	5.742	0.736	11.00
Rumanian Lei	0.1214	0.8022	4.939	−0.162	10.36
Soviet Ruble	0.4059	0.2711	7.483	0.435	43.70

Table 7
GARCH(1,1) with conditional *t*-distribution

	$\hat{\alpha}_1$	$\hat{\beta}_1$	$\hat{v}$	K	S	BL_{15}
Bulgarian Lev	0.0496	0.7857	3.447 (0.620)	3.049	0.230	13.30
Czech Koruna	0.1006	0.7045	5.139 (1.271)	2.994	0.210	14.06
E. German Mark	0.0849	0.8136	3.804 (0.626)	3.072	0.139	15.06
Hungarian Forint	0.2776	0.0398	7.309 (2.946)	2.966	0.027	26.78
Polish Zloty	0.1596	0.6404	4.911 (1.202)	2.879	0.174	23.39
Rumanian Lei	0.1663	0.6334	6.815 (1.996)	3.069	−0.053	10.01
Soviet Ruble	0.1197	0.6000	2.661 (0.483)	2.869	0.086	22.25

However, the estimated degrees of freedom turn out to be extremely low, even falling below 3 for the Soviet Ruble. Though the Forint and the Lei seem to have high values for the degrees of freedom, note that their standard errors are much larger. Furthermore, the symmetric *t* distributions seem to have problems accounting for the asymmetry in some of the currencies, as evinced by the high values of *S* for the Bulgarian Lev and the Czech Koruna. Also, the values of the Box–Ljung statistic indicate that the conditional heteroscedasticity in the Forint, the Zloty and possibly the Ruble has not been accounted for.

We then computed three estimators of the tail index α for all the currencies by the method described in Section 2 above. The three estimators we computed were $\hat{\alpha}_{3,2}, \hat{\alpha}_{3,1}$ and $\hat{\alpha}_{2,1}$ defined by the formula in Eq. (4) above. The value of N_0 was chosen to be 4,000,000 based on the simulations in Deo (2000). For some of the elements of D, the function $\phi_{k,s}$ turned out to be infinite, since $\phi_{k,s}$ involves logarithms. The number of elements for which this happened was negligible and the corresponding terms were dropped from the sum in Eq. (4). The variance covariance matrix of the vector $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}_{3,2}, \hat{\alpha}_{3,1}, \hat{\alpha}_{2,1})'$ was calculated by using the prewhitened kernel estimator described in Deo (2000). The three estimators of α for each of the currencies are reported in Table 8 along with their estimated

Table 8
Estimates of the tail index

	$\hat{\alpha}_{3,2}$	$\hat{\alpha}_{3,1}$	$\hat{\alpha}_{2,1}$
Bulgarian Lev	1.716 (0.056)	1.678 (0.056)	1.660 (0.059)
Czech Koruna	1.860 (0.044)	1.819 (0.051)	1.794 (0.059)
E. German Mark	1.673 (0.068)	1.612 (0.072)	1.583 (0.077)
Hungarian Forint	1.779 (0.130)	1.771 (0.114)	1.762 (0.108)
Polish Zloty	1.745 (0.065)	1.724 (0.071)	1.717 (0.075)
Rumanian Lei	1.867 (0.052)	1.844 (0.061)	1.834 (0.068)
Soviet Ruble	1.566 (0.137)	1.496 (0.113)	1.462 (0.105)

Table 9
z Values and *p* values for goodness-of-fit tests

	Z_1	Z_2	Z_3
Bulgarian Lev	2.044 (0.041)	1.871 (0.061)	1.451 (0.147)
Czech Koruna	2.048 (0.041)	2.044 (0.041)	1.834 (0.067)
E. German Mark	2.638 (0.008)	2.546 (0.011)	2.179 (0.029)
Hungarian Forint	0.325 (0.745)	0.424 (0.672)	0.552 (0.581)
Polish Zloty	0.962 (0.336)	0.810 (0.418)	0.449 (0.653)
Rumanian Lei	1.373 (0.170)	1.202 (0.229)	0.803 (0.422)
Soviet Ruble	2.080 (0.038)	2.089 (0.037)	2.011 (0.044)

standard errors in parentheses. The distribution of the ruble seems to have the thickest tails, with an estimated α around 1.5, while the Rumanian Lei has the thinnest tails with an estimated α of around 1.85.

As before, we computed the three goodness-of-fit tests Z_1, Z_2 and Z_3. Table 9 contains the values obtained for these three test statistics along with the corresponding *p*-values in parentheses. The only currency for which there is very strong evidence against a stable process is the East German mark for which the *p*-values never exceed 0.029. There is some evidence that a stable process is inappropriate for the Soviet Ruble and the Czech Koruna. However, we cannot reject the stable hypothesis for the Bulgarian Lev, the Hungarian Forint, the Rumanian Lei and the Polish Zloty, though the Lev may be a borderline case. Our findings are in contrast with those obtained in Koedijk and Kool (1992) for this same data set. Interestingly enough, they concluded that the stable distribution was inadequate for the Lev, the Koruna, the Forint and the Lei. However, it should be noted that their analysis was on data which was not seasonally adjusted and they did not account for the non-linear dependence in the data in their analysis. This latter fact may also be responsible for their concluding that the data generating process might be changing across different time periods.

Acknowledgements

References

Akgiray, V., Lamoureux, C.G., 1989. Estimation of stable-law parameters: a comparative study. Journal of Business and Economic Statistics 7, 85–93.

Baillie, R., 1993. Comment on modelling asset returns with alternative stable distributions. Econometric Reviews 12, 343–345.

Bloomfield, P., 1976. Fourier Analysis of Time Series: An Introduction. Wiley, New York.

Deo, R., 2000. On estimation and testing goodness-of-fit for m-dependent stable sequences. Journal of Econometrics 99, 349–372.

de Vries, C., 1991. On the relation between GARCH and stable processes. Journal of Econometrics 48, 313–324.

Fama, E., 1963. The behaviour of stock market prices. Journal of Business 38, 34–105.

Hoaglin, D.C., Mosteller, F., Tukey, J.W., 1983. Understanding Robust and Exploratory Data Analysis. Wiley, New York.

Koedijk, K., Kool, C., 1992. Tail estimates of East European exchange rates. Journal of Business and Economic 10, 83–96.

Mandelbrot, B., 1963. The variation of certain speculative prices. Journal of Business 36, 394–419.

Nelson, D., 1990. Stationarity and persistence in the GARCH(1,1) model. Econometric Theory 6, 318–334.

Samorodnitsky, G., Taqqu, M., 1994. Stable Non-Gaussian Random Processes. Chapman & Hall, New York.