Centre for Computational Finance and Economic Agents

Working Paper Series

WP008-06

Peter Winker
Vahidin Jeleskovic

The Unconditional Distribution of Exchange Rate Returns: Statistics, Robustness, Time Aggregation

October 2006

University of Essex

CCFEA

www.ccfea.net

The Unconditional Distribution of Exchange Rate Returns: Statistics, Robustness, Time Aggregation*

Peter Winker[†] Vahidin Jeleskovic[‡]

October 4, 2006

Abstract

Empirical research on financial market data reports a number of properties of return time series which are considered as 'stylized facts'. In this contribution, statistics for the unconditional distribution of foreign exchange rate returns are discussed. The robustness of these properties is assessed using the bootstrap and considering different subsamples. The effect of time aggregation is also analyzed.

Based on the results robust stylized facts are selected which – together with similar results for conditional distributions and long range dependence provided in a companion paper – will result in an objective function for an indirect estimation method of agent based models.

Keywords: Excess kurtosis; Jarque-Bera statistic; Kolmogorov-Smirnov statistic; fat tails; Pareto law; tail index.
JEL-classification: C14, C15, F31

*Research has been supported by the DFG grant WI 20024/2-1/2. We are indebted to M. Gilli for helpful comments on a preliminary draft of the paper and to J. Brüggemann for valuable research assistance.

[†]Department of Economics, University of Giessen
[‡]Department of Economics, University of Giessen

1 Introduction

Although market efficiency is accepted as a general guideline for the analysis of financial markets (Fama, 1970; Fama, 1991), there is also a general consensus that financial market time series are not well described by a pure random walk with normal and independent innovations. In fact, 'stylized facts' such as fat tails, volatility clustering and long memory are often mentioned in the literature (Cont, 2001). Such stylized facts are used as a base for extending the theoretical understanding of markets and to generate new models.

In fact, taking into account bounded rationality and heterogeneity of market participants (Shleifer, 2000), agent based models of financial markets appear to be able to match some of the stylized facts of foreign exchange markets.[1] In fact, this property of agent based models is often used as argument in favour of this class of models. For example, Lux (1998) and Tay and Linn (2001) use the first four centered moments, while Arthur *et al.* (1999) use the standard deviation and the kurtosis as one of the main benchmarks for their agent based models. Furthermore, Lux and Marchesi (2000) also consider the tail index.[2] When using these characteristics of exchange rates as a benchmark for agent based models, one has to take into account the robustness of these statistics, i.e. the extent to which their calculated values depend on a specific (small) sample under consideration. Consequently, identifying robust properties of the data might eventually help to estimate parameters of agent based models and to test the validity of the fitted model. However, methods for this type of econometric validation are still in their infancy (Gilli and Winker (2003) and Alfarano *et al.* (2005)).

In this paper, we restrict our analysis to the foreign exchange market.[3] As a base for a further assessment of agent based models, we start by determining which of the 'stylized facts' are actually stylized facts in the sense of nontrivial statistical properties, which are robust over time and across assets (Cont, 2001). Furthermore, we measure distributional properties of these stylized facts using the bootstrap approach.

---

[1] E.g., Lux and Schornstein (2005) use a set of characteristics as empirical benchmarks for a genetic learning based model of foreign exchange markets.

[2] Properties of the conditional return distribution are also considered in this context. These properties are discussed in the companion paper Winker and Jeleskovic (2006, forthcomming).

[3] See de Vries (1994), Cont (2001) and Krämer (2002) for summaries of statistical properties of financial market data.

Given the huge number of statistical properties of foreign exchange data mentioned in the literature, we concentrate on the unconditional return distribution in this paper. A companion paper deals with conditional distributions and long range dependence (Winker and Jeleskovic, 2006, forthcomming). Of course, the set of properties covered in this contribution cannot be exhausting. Nevertheless, it should be ample enough to select appropriate statistics for the purpose of indirect estimation of agent based models as introduced by Gilli and Winker (2003).

We proceed as follows. Section 2 introduces the data base used for the empirical analysis and their empirical moments. The shape of the univariate return distribution, in particular the fat tail property, is analyzed in Section 3 and suitable statistics are presented. The robustness of the considered statistics is assessed in Section 4 by analyzing subsamples of the data and application of the bootstrap method. Finally, the impact of time aggregation is covered in Section 5. Section 6 summarizes our main findings and the implications for indirect estimation approaches.

2 Data and Empirical Moments

2.1 The Data

For the empirical analysis, we use foreign exchange time series at daily frequency (5-day-week). The sample includes 8088 observations of the German D-Mark (DM), British Pound (GBP), French Franc (FRF), Swiss Franc (CHF), and the Japanese Yen (JPY) against the US Dollar in volume notation for the period 1st of January 1974 to 4th of January 2005.[4] The values for the German D-Mark and the French Franc are the implicit rates after the introduction of the Euro in 1999. Each time series contains 284 missing values due to bank holidays. We use the sample 1.1.1975 to 31.12.2004 for the following analysis if not stated otherwise. For the observations following a missing value, the return is calculated on the base of the last preceding valid observation. The return time series are calculated as logarithmic differences of the price series and denoted by r_i, $i \in \{DM, GBP, FRF, CHF, JPY\}$.

Figure 1 shows the exchange rate time series for the German D-Mark against the US Dollar in the upper panel, and the corresponding return time

[4]The data were retrieved from http://fx.sauder.ubc.ca/data.html, a site of the Sauder School of Business, University of British Columbia.

series in the lower panel.

Figure 1: Exchange Rate D-Mark against US Dollar and Returns

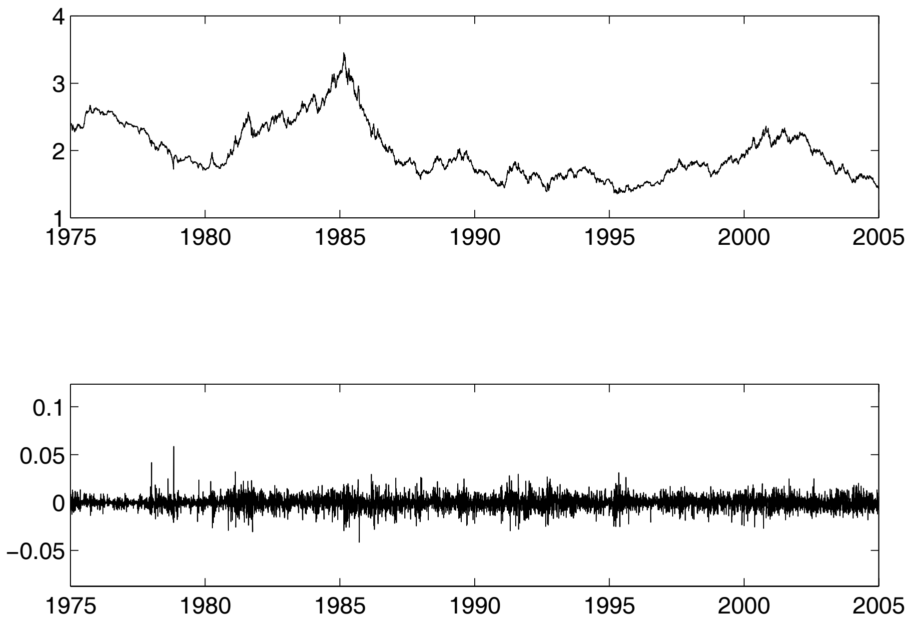

2.2 Empirical Moments and Tests for Normality

In Table 1, we start by providing the empirical skewness, kurtosis and two test statistics for normality of the distribution (Jarque-Bera and Kolmogorov-Smirnov). While the empirical skewness is close to zero for most series, the observed kurtosis is clearly above the value of three for a normal distribution resulting in highly significant Jarque-Bera (JB) statistics.[5] Although some measures of tail behaviour discussed in the next section challenge the existence of higher order moments, these results strongly support the rejection of the normal assumption. This finding is supported by the values of the Kolmogorov-Smirnov statistic (KS) reported in the last column. For the test of the null hypothesis of a normal distribution, the moments of the normal distribution have to be estimated. Therefore, the modified statistic and

[5]Under the null hypothesis of a normal distribution, the Jarque-Bera statistic follows asymptotically a χ^2-distribution with two degrees of freedom.

critical values from Lilliefors (1967) are used for the test. Again, the null hypothesis has to be rejected at any conventional level of significance.

Table 1: Test Statistics for Normality Assumption

Series	Obs.	Mean	Std.dev.	Skew.	Kurt.	JB	KS
r_{DM}	7554	-6.8E-05	0.0065	0.0314	5.987	2809	0.0490
r_{FRF}	7554	1.1E-05	0.0064	0.1000	7.169	5483	0.0544
r_{GBP}	7554	2.6E-05	0.0061	0.1300	6.677	4276	0.0633
r_{CHF}	7554	-0.00011	0.0076	0.0438	13.699	36033	0.0527
r_{JPY}	7554	-0.00014	0.0066	-0.4858	7.377	6328	0.0704

It should be kept in mind that the goal of our analysis does not consist in estimating structural parameters of some theoretical stochastic process generating price or return series.[6] Instead, we aim at identifying robust features of financial market time series which can be used for estimating the parameters of agent based models. Obviously, a statistic like the Kolmogorov-Smirnov statistic, which is based on the whole distribution of returns, might provide more information than a single moment estimator, in particular, when the distributions are clearly not normal.

3 Fat Tails

3.1 Stable Distributions, Extreme Value Theory and Tail Index

From the results presented in the previous section, it becomes obvious that normal distributions are not well suited to model the unconditional distribution of returns. Several alternative parametric models have been proposed to match the stylized fact of heavy tails.[7] One alternative distributional assumption consists of the *stable* class.[8] The stable distributions are a gener-

[6]Consequently, the requirement of ergodicity (Cont, 2001), e.g., is not necessarily binding for our application.

[7]This effect becomes less pronounced under time aggregation (see Section 5).

[8]This class of distributions was characterized by Lévy (1924) in his study of the normalized sums of *iid* terms.

alization of the normal distribution. They are characterized by the condition that a sum of stable random variables is also a stable random variable. The general stable distribution can be described by four parameters: an index of stability $\alpha \in (0, 2]$ (*tail index*), a skewness parameter, a scale parameter and a location parameter.

For values of the tail index $\alpha < 2$, the stable distribution becomes so heavy tailed that moments of order two and more do not exist. Consequently, sample estimates of, e.g., variance and kurtosis will not converge as the sample size increases. Infinite higher moments do not only result in theoretical modelling problems, but also suggest that empirical variance and higher moments should increase as the sample size increases. This result is not supported by empirical findings. Finally, stable distributions do not allow to explain the empirical finding that return distributions become closer to normal distributions under time aggregation (see Section 5). These difficulties led to the development of models based on a mixture of distributions and conditional distributions, e.g., GARCH-models (Winker and Jeleskovic, 2006, forthcomming).

Instead of searching for a parametric model of the return distribution, an alternative approach concentrates on extreme events, which correspond to the tails of the distribution. Extreme value theory (EVT) is a formal framework to study this tail behavior, in particular of fat-tailed distributions. The EVT provides models of the asymptotic characteristics at the tails of a distribution, e.g. of stationary return series. Often this approach provides a better fit to extreme quantiles than the conventional parametric approaches (Embrechts, 2000a; Gençay *et al.*, 2003a). Furthermore, the EVT allows to make inferences about the return distribution beyond the observed range of sample returns. Similarly to the normal distribution which is the limiting distribution for sample sums in a central limit theorem, the family of extreme value distributions is used to study the limiting distributions of the sample maxima.

To introduce the basic notion of the EVT, let us suppose that a sequence r_t, $t = 1, \ldots, T$ of *iid* random variables[9] is given with unknown cumulative distribution function F. We denote the corresponding order statistics by $r_{1,T} \leq r_{2,T} \leq \ldots \leq r_{T,T}$ and assume that the properly centered and nor-

[9] The assumption of independent random variables can be dropped and the theoretical results still hold (McNeil, 1997; Gençay *et al.*, 2003b). The assumption of identical distribution can also be relaxed (Gençay *et al.*, 2003b).

malized sample maxima $r_{T,T} = \max\{r_1, r_2, \ldots, r_T\}$ for $T \to \infty$ converge in distribution to a non-degenerating limiting distribution function H. Then, for a given sequence of $a_T > 0$ and b_T, $(r_{T,T} - b_T)/a_T$ necessarily converge to the non-degenerate limit H, which is called generalized extreme value (GEV) distribution (von Mises, 1954; Jenkinson, 1955):

$$H_\xi(r) = \begin{cases} e^{-(1+\xi r)^{-1/\xi}} & \text{if} \quad \xi \neq 0 \\ e^{-e^{-r}} & \text{if} \quad \xi = 0\,, \end{cases} \tag{1}$$

where $1 + \xi r > 0$ and the shape parameter $\xi \in \mathbb{R}$. The tail index is defined as $\alpha = \xi^{-1}$.

For $\xi = 0$, the distribution has exponentially decreasing tails (the Gumbel class) and for $\xi < 0$ one obtains the Weibull class with short-tailed distributions. If $\xi > 0$, the distribution is called Fréchet-type (or fat-tailed) and the tail decays with a power law. The larger the shape parameter, the more fat-tailed the distribution. Thus, an estimate of ξ or α, respectively, provides a measure of the fat-taildness of the distribution. In particular, the theoretical p-th moment of the return distribution is only existent if $p < \alpha$ (Embrechts *et al.*, 1999, p. 165).

3.2 Estimation of the Tail Index

The most popular estimator for the tail index of Pareto-type distributions is the conditional maximum likelihood estimator introduced by Hill (1975).[10] The Hill estimator is based on the order statistics exceeding a high threshold value $m > 0$, which is assumed to be known.[11] Let $k = k(T)$ denote the number of order statistics exceeding the threshold. Then, the maximum likelihood estimator of ξ conditional on m suggested by Hill is given by

$$\hat{\xi}_k = \frac{1}{k} \sum_{j=1}^{k} log(r_{T-j+1,T}) - log(r_{T-k,T}) \tag{2}$$

In order to obtain satisfactory properties of the Hill estimator, the threshold m has to be chosen such that the following conditions hold (Wagner and Marsh, 2003):

$$k(T) \to \infty, \qquad k(T)/T \to 0, \qquad \text{as} \qquad T \to \infty \tag{3}$$

[10] A comparison with alternative estimators is provided by Kearns and Pagan (1997).

[11] For the lower tail, a symmetric construction is used.

If these conditions are fulfilled, the Hill estimator exhibits the following properties (Wagner and Marsh, 2003):

1.) The estimator is consistent (Embrechts *et al.*, 1999),

2.) under some additional assumptions on the tail behavior of F, the estimator is asymptotic normally distributed with some asymptotic bias term (de Haan and Peng, 1998; Segers, 2001),

3.) a theoretically derived optimal rate of convergence is achieved, and

4.) the estimator is asymptotically quite robust with respect to deviations from independence (Hsing, 1991), e.g. in a GARCH sense (Resnick and Starica, 1998).

In Table 2, we provide the results for the estimation of the tail index $\alpha = 1/\xi$ using three different values for the threshold value m, the 2.5%, the 5% and the 10% quantile of the empirical distribution. The first three columns provide the estimates for the left tail, while the last three columns show the corresponding results for the right tail of the return distributions.

Table 2: Hill Estimator of Left and Right Tail Index α

	Left Tail			Right Tail		
Threshold	2.5%	5%	10%	2.5%	5%	10%
r_{DM}	4.69	3.98	3.45	4.29	3.98	3.28
r_{FRF}	4.38	3.56	3.27	4.10	3.73	3.12
r_{GBP}	3.74	3.77	3.40	3.81	3.68	2.97
r_{CHF}	4.52	3.69	3.24	3.97	3.80	3.39
r_{JPY}	3.39	3.38	2.74	4.63	3.69	3.10

The values ranging from 3.3 to 4.0 for a 5% threshold are in line with usual estimates for foreign exchange data (Lux and Schornstein, 2005, p. 177).[12] Nevertheless, comparing the column entries, it becomes evident that a critical aspect of the Hill estimator is the choice of the threshold value m. One way to deal with the selection of a suitable threshold or sample fraction $k(T)$

[12]Similar values are also obtained by Gilli and Këllezi (2006) for several financial market indices including stock prices (see also Jansen and de Vries (1991)).

consists in analyzing the so-called *Hill plot*, which is based on a series of Hill estimates (Drees, 1998):

$$\{(k, \hat{\xi}_{k,T}) : 1 \leq k \leq T\} \tag{4}$$

The tail index estimator should be chosen from a stable region of the Hill plot (Gilli and Këllezi, 2006). Figure 7 provides the Hill plot for the right tail of r_{DM} (the results for the other series can be found in the appendix).

Figure 2: Hill Plot for r_{DM}

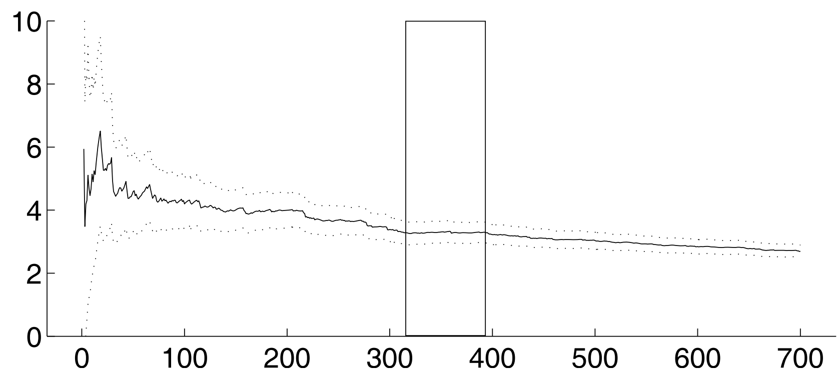

Obviously, the Hill estimator (solid lines) and its 95% confidence bands (dashed lines) exhibit marked fluctuations for very small values of k. The rectangles indicate stable regions for the different exchange rate data. This stable regions are from 8% to 10% for r_{DM}, 7% to 10.4% for r_{FRF}, 9.4% to 10.5% for r_{GBP}, 8.5% to 9.8% for r_{CHF} and 6% to 7.1% for r_{JPY}. Thus, the standard choice of the threshold value of 5% or 10% seems adequate for our data. In Table 3, we provide the average values of the Hill estimator over these stable regions as well as for the region 5% to 10%.[13]

Summarizing the results of the estimation of left and right tail indices, there is overwhelming evidence for a tail index smaller than four, i.e. heavy tailed return series. Consequently, although the empirical kurtosis might still be useful as empirical moment, it does not provide an estimator of the (non

[13]The stable regions for the left tail are qualitatively similar to those for the right tail. The exact values are available on request.

Table 3: Average Hill Estimator of Left and Right Tail Index α

Series	r_{DM}	r_{FRF}	r_{GBP}	r_{CHF}	r_{JPY}
	Left Tail				
Average α (graphical)	3.35	3.41	3.60	3.40	3.41
Average α (5%-10%)	3.54	3.29	3.26	3.50	3.39
	Right Tail				
Series	r_{DM}	r_{FRF}	r_{GBP}	r_{CHF}	r_{JPY}
Average α (graphical)	3.45	3.90	3.76	3.41	3.50
Average α (5%-10%)	3.62	3.40	3.61	3.47	3.10

existent) theoretical fourth moment of the return series. The same reasoning applies to the Jarque-Bera statistic.

3.3 BM- and POT-Method

One shortcoming of the Hill-estimator lies in the fact that the approximation in (1) relies on unknown scaling parameters. A different approach is the use of the three parameter representation of the GEV, the so-called generalized Pareto distribution (GPD), for r based on $\frac{r-\mu}{\sigma}$, where μ is the location and σ is the scale parameter (von Mises, 1954). In order to obtain estimates of these parameters, two methods are used.

First, the *Method of Block Maxima (BMM)* is based on observed extreme observations in subsamples of the data. To this end, the data set is divided in n blocks of equal length and the maximum value for each block is calculated and stored. The GPD is fitted to this set of block maxima by means of maximum likelihood.[14] For the choice of the block length we closely follow Gilli and Këllezi (2006) using the calendar year. However, in order to allow for the bootstrap analysis, we fix the block length to 250 days. The estimates of the tail indices obtained with this method are shown in Table 4. These estimates are highly sensitive with regard to the chosen block length.

Second, the *Peak-over-Threshold Method (POT)* uses available data more efficiently for estimating the tail index. This method is based on extracting extreme observations from the empirical distribution function $F(x) = \mathrm{P}(r_t \leq$

[14]For the estimation of the GPD to the left tail the returns are multiplied with -1.

x) for a given high threshold value m. The excess over m is defined by $z = r - m$, which is also called the peak over threshold for the returns $r > m$. Then, the distribution of excess returns conditional on exceeding the threshold is given by

$$F_m(z) = \frac{\mathrm{P}(r - m \leq z | r > m)}{\mathrm{P}(r > m)} = \frac{F(z+m) - F(m)}{1 - F(m)}\,, \tag{5}$$

and therefore

$$F(r) = [1 - F(m)]F_m(z) + F(m) \;\text{ with } z = r - m\,, \tag{6}$$

for $r > m$. For increasing thresholds m, the excess distribution $F_m(z)$ converges to the GPD.[15] $F(m)$ can be estimated by $(T - T_m)/T$, where T_m is the number of returns r_t exceeding the threshold m. For given threshold m, setting $\mu = m$ allows to obtain maximum likelihood estimates for σ and ξ. As before, the estimate of the tail index is given by $\alpha = \xi^{-1}$.

The results for the POT method depend crucially on the choice of m. The convergence result requires to choose a high value for m. However, large values of m will result in a small number of excess returns and a poor approximation of (6) for a given sample size T. The sample mean excess function (MEF) provides a means for selecting the threshold m. The MEF is defined by

$$e_T(m') = \frac{\sum I_{(r>m')}(r - m')}{\sum I_{(r>m')}}\,, \tag{7}$$

where I is the indicator function. The threshold m will be chosen such that the portion of the MEF is positiv and linear for $m' > m$. Figure 3 shows plots of the MEF for the left (upper panel) and right (lower panel) tail of r_{DM}.

While it appears feasible to select m by visual inspection of the MEF plot for empirical data, this approach cannot be followed for the bootstrap analysis. Therefore, an automatic procedure for detecting the threshold m has to be used. To this end, linear regressions of $e_T(m')$ against m' are run for $m = T - 150, \ldots, T - 10$. For each regression, the Schwarz information criterion (BIC) is calculated. Then, the threshold corresponding to the smallest value of BIC is considered. If this value coincides with the lower

[15] See the references provided in Gilli and Këllezi (2006) to work by Balkema and de Haan(1974) and Pickands (1975).

Figure 3: Mean excess function for r_{DM}

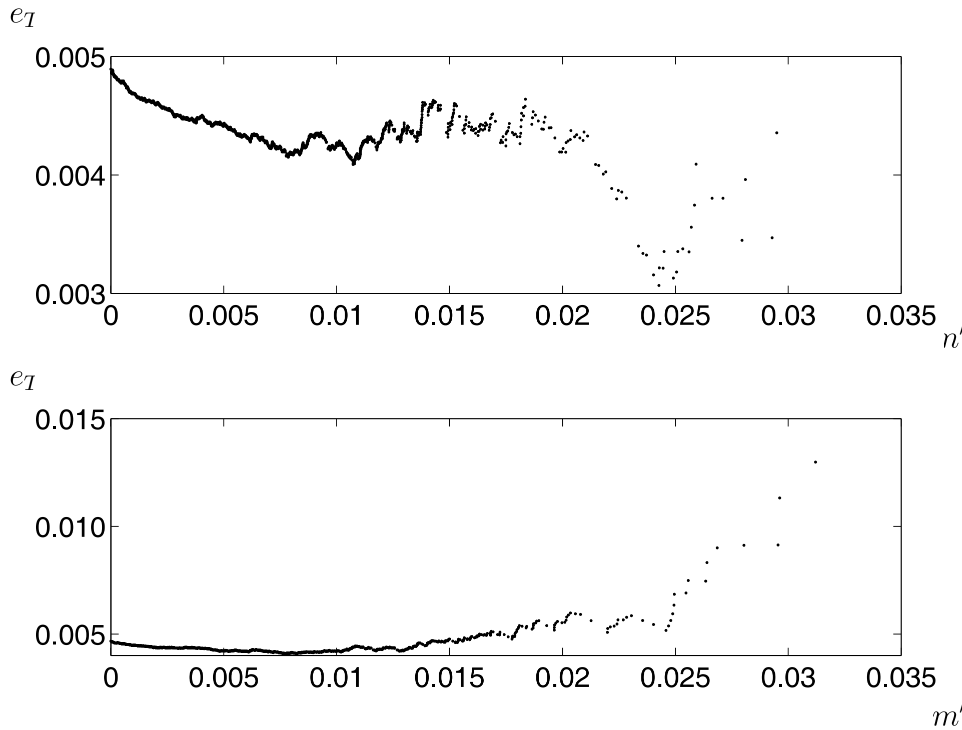

bound ($T - 150$), it is increased in steps of size 5 until a local minimum is found.[16] For the results presented in Table 4, the threshold m selected by this approach is indicated in parentheses. The same procedure is also used for each bootstrap replication in Subsection 4.2.

It should be noted that the results for the POT method are highly sensitive with regard to the choice of m. This sensitivity might explain the high dispersion of the estimates provided in Table 4. Although a sample size of 7554 is typically regarded as quite sensible, it appears too small to obtain reliable estimates of the tail index based on the BMM and POT methods given the high threshold values m identified both based on visual inspection of the MEF plot and the automatic procedure described above.[17] This

[16]Another option is to take a fixed portion of the data length (McNeil and Frey, 2000).

[17]Gilli and Këllezi (2006) and Gençay and Selçuk (2004) report better results for stock market time series. These time series exhibit more extreme observations. Thus, the Pareto distribution can be fitted based on a larger number of observations than for the exchange

finding is also supported by the results of the bootstrap analysis provided in Subsection 4.2. Thus, it is questionable, whether the tail index estimates obtained by the BMM and POT methods can be used as a benchmark for the analysis of models of exchange rate time series.

Table 4: BBM and POT Estimators of Left and Right Tail Index α

Method	Tail	r_{DM}	r_{FRF}	r_{GBP}	r_{CHF}	r_{JPY}
BMM	α_{left}	1.6264	1.5607	1.5650	1.7265	1.5720
BMM	α_{right}	1.6234	1.5750	1.5608	3.1652	1.5608
POT	α_{left}	4.0508	5.4496	3.7721	3.3183	3.0420
		(0.0248)	(0.0252)	(0.0260001)	(0.0222)	(0.02697)
POT	α_{right}	3.8590	3.0335	9.6023	2.1223	7.5520
		(0.0231)	(0.0238)	(0.0280)	(0.0243)	(0.01614)

4 Robustness

According to the empirical findings presented in the previous section, the unconditional distribution of exchange rate returns is clearly non normal. In fact, both the empirical kurtosis and the Hill-estimator of the tail index indicate heavy tails for all exchange rates considered. However, in order to be considered as a 'stylized fact', which can be used as a base for further analysis, i.e. indirect estimation of the parameters of agent based models, it is of importance how robust these findings are.

In order to assess the robustness of the results, we employ two methods. First, the calculations are repeated for different subsamples of the data in order to allow for a comparison of the estimates over time. Thereby, both a rolling window analysis and the analysis for long subsamples are considered. Second, the distribution of the statistics is estimated by a bootstrap implementation.

---

rate time series.

4.1 Subsamples

We start with a rolling windows analysis considering overlapping subsamples of length 200 and 400 days, respectively. The lower moments of the unconditional returns exhibit marked fluctuations over time when considering these small subsamples. However, neither for the means nor for the standard deviations of the exchange rate returns a clear trend can be discovered. Figure 4 plots the results for the mean returns over overlapping 400 days windows as an example. A similar picture emerges for overlapping windows of length 200.

Figure 4: Mean of Exchange Rate Returns (in percent) for Rolling Windows (400 Days)

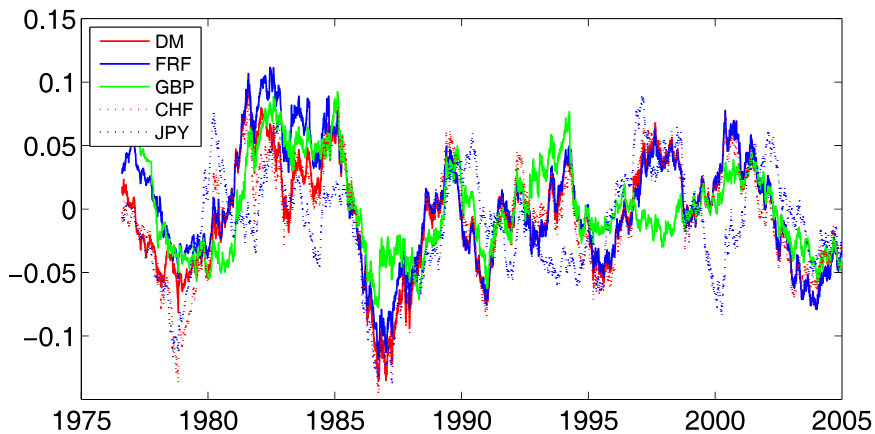

The results for the empirical skewness and empirical kurtosis are less robust. While the skewness exhibits a few extreme values during the first 1500 days (2500 days for r_{CHF}), it sticks to values close to zero for the rest of the sample. Similarly, as shown in Figure 5, the empirical kurtosis reaches very high values during the first 1500 days (2500 days for r_{CHF}).[18] The dashed horizontal line corresponds to the theoretical kurtosis of a normal distribution. Thus, although the extreme values during the first sample years

[18]Although the extreme values reach 75, the plot is censored at 20 in order to provide more information on the stable period.

indicate that the empirical kurtosis might not be a robust measure of fat-taildness of the distribution,[19] at least excess kurtosis appears to be a robust finding. Again, the results using a window length of 200 days are qualitatively similar.

Figure 5: Kurtosis of Exchange Rate Returns for Rolling Windows (400 Days)

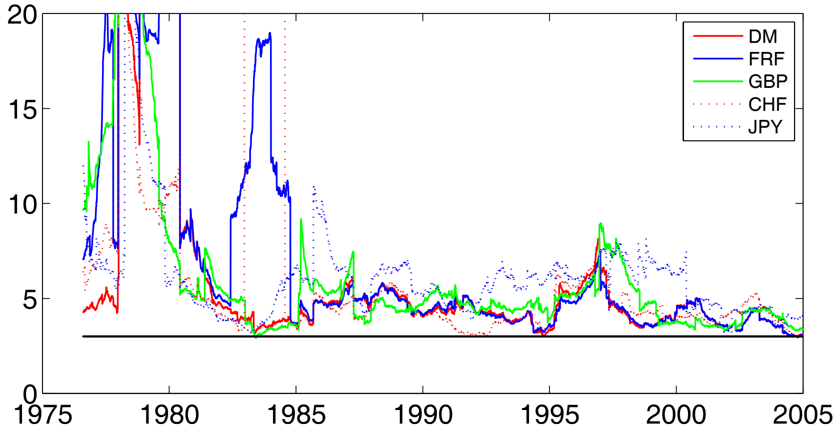

Given the missing robustness of the empirical kurtosis during the first sample years, a consideration of the tail index would be of interest. However, for obvious reasons the rolling window analysis cannot be extended to the Hill-estimator of the tail index.

However, the results suggest to consider different subperiods of the sample. For this purpose, we consider the periods 1975-1984, 1985-1994, 1995-2004 separately. For this ten year periods, it is also possible to obtain estimates of the tail-index. Table 5 summarizes our findings for the three subperiods. These results are in line with the findings for the rolling windows analysis. While for the first two moments no clear tendency can be observed, skewness tends to become negative after the first subperiod. Furthermore, the return distributions tend to become less fat tailed. Nevertheless, they are still far from being normal for all three subperiods.

[19] This result might have been expected based on the estimates of the tail index as discussed in the previous section.

Table 5: Test Statistics for Subperiods

Series	Obs.	Mean	Std. dev.	Skew.	Kurt.	JB	KS	Tail Index (5% - 10%) Left	Right
1975 - 1984									
r_{DM}	2505	0.0001	0.0059	0.337	9.855	4953	0.075	3.474	3.370
r_{FRF}	2505	0.0003	0.0058	0.442	7.629	2318	0.087	2.634	2.830
r_{GBP}	2505	0.0003	0.0060	0.617	14.037	12874	0.093	3.569	3.282
r_{CHF}	2505	4E-06	0.0078	0.363	28.872	69920	0.089	3.246	2.911
r_{JPY}	2505	-7E-05	0.0059	-0.579	8.655	3477	0.101	3.466	3.451
1985 - 1994									
r_{DM}	2532	-0.0003	0.0072	-0.035	4.884	375	0.052	3.637	3.454
r_{FRF}	2532	-0.0001	0.0073	-0.011	5.789	820	0.058	3.526	3.494
r_{GBP}	2532	-0.0002	0.0070	-0.043	4.916	388	0.052	3.132	3.823
r_{CHF}	2532	-0.0003	0.0079	-0.070	4.325	187	0.038	3.954	3.897
r_{JPY}	2532	-0.0004	0.0064	-0.327	5.937	955	0.056	3.088	3.344
1995 - 2004									
r_{DM}	2517	-3E-05	0.0063	-0.078	4.297	179	0.037	4.010	4.092
r_{FRF}	2517	-8E-05	0.0050	0.121	4.439	223	0.042	4.323	4.026
r_{GBP}	2517	-4E-05	0.0062	-0.137	4.080	130	0.035	5.138	4.051
r_{CHF}	2517	-6E-05	0.0069	-0.232	4.689	322	0.040	3.516	4.227
r_{JPY}	2517	1E-05	0.0074	-0.545	7.226	1998	0.059	2.923	3.434

A formal test for constancy of the tail index is introduced by Jansen and de Vries (1991, p. 20).[20] It cannot be applied directly to the average estimates of the tail index provided in Table 5, but requires a fixed threshold m for each estimate. Therefore, we repeated the analysis for a fixed threshold of $m = 100$ falling in the interval 5%-10%. When comparing the current subperiod with the previous one using the estimate of the previous subperiod as benchmark, we find a significant change (at the 5% level) of the tail index for the following cases. For the left tail, r_{FRF} between the first and second subperiod, and r_{GBP} between the second and third subperiod. For the right

[20] An alternative approach to test for structural changes in the tail behaviour is proposed by Quintos *et al.* (2001).

tail, no significant change has been found.[21]

4.2 Bootstrap Analysis

A generalization of the subsample approach is the general bootstrap, i.e. a general resampling procedure from the empirical return distributions. Although only the unconditional distribution is considered in this paper, we use both a simple bootstrap and the block bootstrap with differing block length. In particular, the block bootstrap is used for comparability with the results for conditional distributions presented in Winker and Jeleskovic (2006, forthcomming). However, it is expected that the results for the unconditional return distributions do not differ between simple and block bootstrap.

In Figure 6, the results of the simple bootstrap are plotted by means of normal kernel plots of the distribution obtained for 1 000 bootstrap drawings.[22] In addition to the kernel estimates, a normal (lognormal for the JB-statistic) approximation to the data is shown as dotted lines in the plots.

Table 6 shows summary statistics for the distribution of r_{DM} obtained by the simple bootstrap. The corresponding results for the other return distributions are provided in the appendix. As expected for statistics of the unconditional distributions, the results for the block bootstrap with window length 20 days do not differ substantially.[23]

The results of the bootstrap support our findings from the subsample analysis. The mean of foreign exchange rate returns is not a stable characteristic as it measures basically the trend of the exchange rate over the whole sample considered. Thus, the large bootstrap variance of the mean is not surprising. By contrast, the bootstrap values of the unconditional variance are surprisingly stable. However, we have also to consider conditional heteroskedasticity (see Winker and Jeleskovic (2006, forthcomming)). The skewness exhibits large fluctuations in the bootstrap distribution. This might be a result of the finding that skewness is not significantly different from zero for the whole sample. Thus, the bootstrap distribution might correspond to random fluctuations. For the empirical kurtosis, the bootstrap

[21] The alternative methods for estimating the tail index (BMM and POT) are not considered on subsamples given the difficulties to obtain reasonable estimates for the whole sample.

[22] The corresponding results for a block bootstrap with a block length of 20 days are shown in Figure 8 in the appendix.

[23] These results are available on request from the authors.

Figure 6: Bootstrap Distributions (Simple Bootstrap)

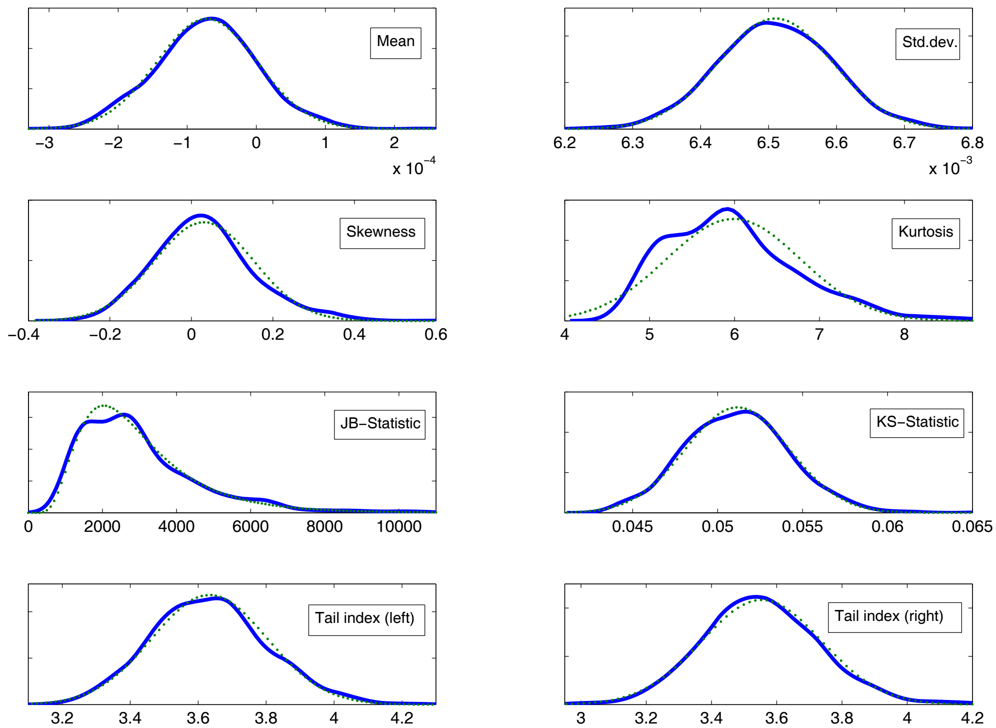

standard deviation is smaller than for the skewness in relative terms. Nevertheless, both for the empirical kurtosis and the Jarque-Bera statistic we have to keep in mind that the corresponding theoretical moments might not exist. In fact, the bootstrap results for the tail index estimates (Hill estimator) are quite robust indicating that the tail index is significantly smaller than 4 for all return time series (see also the results provided in the appendix). However, the estimates of the tail index obtained from the BMM and POT methods are very unstable. This finding is due to their dependency on a very small number of extremum observations. Finally, the bootstrap estimates of the Kolmogorov-Smirnov statistic fall in a rather small intervall. Thus, this statistic might also be considered as a robust characteristic of the data.

Table 6: Bootstrap Distribution of r_{DM}

Statistic	DM/US	Mean	5%	50%	95%	Std.dev.
Mean	-0.68E-4	-0.71E-4	-1.98E-4	-0.72E-4	0.49E-4	0.75E-4
Std.dev.	0.0065	0.0065	0.0064	0.0065	0.0066	0.0001
Skewness	0.0314	0.0255	-0.1592	0.0165	0.2241	0.1217
Kurtosis	5.9866	5.9503	4.9149	5.8387	7.4826	0.7985
JB-Statistic	2806	2957	1175	2534	6351	1673
KS-Statistic	0.0490	0.0513	0.0465	0.0512	0.0566	0.0031
Tail index (left)	3.6208	3.6300	3.3709	3.6272	3.9091	0.1656
Tail index (right)	3.5409	3.5602	3.2598	3.5607	3.8545	0.1809
BMM α_{left}	1.6264	-0.6339	-37.4834	-2.0654	36.7693	177.4995
BMM α_{right}	1.6234	2.8987	-16.8241	2.8594	25.8162	66.1550
POT α_{left}	4.0508	4.4393	-32.7288	-3.1778	28.0285	165.2139
POT α_{right}	3.8590	3.7838	-20.7924	2.6654	19.6530	59.0284

5 Time aggregation

The final step of our analysis points towards the effects of time aggregation on the statistical properties of unconditional exchange rate return distributions. Obviously, the mean return will not be affected by considering the observations at lower frequency, i.e. weekly or monthly returns.[24] Nevertheless, other statistics might be affected. In particular, it is often reported that financial market returns become closer to a normal distribution at lower frequencies.

Therefore, we recalculate all statistics presented in Tables 1 and 2 for blocks of 5, 20 and 60 observations corresponding roughly to weekly, monthly and quarterly data, respectively. Table 7 shows the results for the standard test statistics including the results for daily returns from Table 1 as a benchmark.

In contrast to the mean, the standard deviation shows a clear increasing trend at a rate of close to square root of the interval length. If the returns were independently distributed, this is the expected result. Although Winker

[24] By forming non overlapping intervals, some observations might get lost resulting in small changes of the overall sample and corresponding effects on the statistics.

Table 7: Test Statistics under Time Aggregation

Series	Obs.	Mean	Std.dev.	Skew.	Kurt.	JB	KS
1-day returns							
r_{DM}	7554	-6.8E-05	0.0065	0.0314	5.987	2809	0.0490
r_{FRF}	7554	1.1E-05	0.0064	0.1000	7.169	5483	0.0544
r_{GBP}	7554	2.6E-05	0.0061	0.1300	6.677	4276	0.0633
r_{CHF}	7554	-0.00011	0.0076	0.0438	13.699	36033	0.0527
r_{JPY}	7554	-0.00014	0.0066	-0.4858	7.377	6328	0.0704
5-day returns							
r_{DM}	1510	-0.00034	0.0152	0.0642	5.605	428	0.0424
r_{FRF}	1510	0.00013	0.0145	0.3083	6.806	936	0.0503
r_{GBP}	1510	5.8E-05	0.0149	0.1377	5.710	467	0.0528
r_{CHF}	1510	-0.00054	0.0169	-0.0564	5.443	376	0.0522
r_{JPY}	1510	-0.00071	0.0154	-0.9211	9.873	3186	0.0736
20-day returns							
r_{DM}	376	-0.00123	0.0323	-0.2210	3.147	3.402	0.0349
r_{FRF}	376	0.00055	0.0297	0.1362	4.086	19.640	0.0399
r_{GBP}	376	0.00035	0.0314	-0.0279	3.179	0.553	0.0454
r_{CHF}	376	-0.00206	0.0353	-0.1748	3.599	7.536	0.0431
r_{JPY}	376	-0.00278	0.0333	-0.4327	4.063	29.429	0.0643
60-day returns							
r_{DM}	125	-0.00333	0.0585	0.1476	2.640	1.130	0.0700
r_{FRF}	125	0.00212	0.0579	0.5731	5.426	37.494	0.0446
r_{GBP}	125	0.00146	0.0576	0.2674	2.348	3.702	0.0580
r_{CHF}	125	-0.00571	0.0652	0.0527	2.593	0.919	0.0488
r_{JPY}	125	-0.00806	0.0620	-0.2355	4.106	7.525	0.0535

and Jeleskovic (2006, forthcomming) show that the returns considered in this application exhibit some time dependence (like GARCH-effects or long-range dependence), the convergence of the standard deviation of mean returns under time aggregation does not seem to be affected to a relevant extent. The repeated change in sign of the skewness is more difficult to explain. For the kurtosis we find a general decreasing trend with increasing time intervals. Consequently, while using the JB-statistic, normality of returns has to be rejected for all return time series at a 5% level for daily and weekly data, it can be rejected only for r_{FRF}, r_{CHF}, and r_{JPY} for a 20-day horizon and for r_{FRF} and r_{JPY} at a 60-day horizon. Using the KS-statistic in the Lilliefors-test, the null hypothesis of normally distributed returns cannot be rejected for any exchange rate at the 60-day horizon.

These findings for a tendency to normality of exchange rate returns under time aggregation are supported by the results for the estimators of the tail indices presented in Table 8. These estimators are calculated as the average of the Hll-estimators over the interval 5%-10%. Although the estimates for period lengths of 20 and, in particular, 60 days become less reliable due to the small sample size, the overall tendency to increasing values of the tail index is obvious. In particular, for a period length of 20 days, all return time series exhibit a tail index larger than four corresponding to the existence of the theoretical fourth moment (kurtosis).

Table 8: Hill-estimator under temporary aggregation.

Period	1 day		5 days		20 days		60 days	
Tail Index	Left	Right	Left	Right	Left	Right	Left	Right
r_{DM}	3.54	3.62	3.91	3.79	4.71	7.43	5.01	23.37
r_{FRF}	3.29	3.40	3.91	3.60	6.39	5.79	12.35	36.64
r_{GBP}	3.26	3.61	3.36	3.19	4.10	5.03	6.33	3.05
r_{CHF}	3.50	3.47	4.00	3.99	5.08	5.04	6.73	23.97
r_{JPY}	3.39	3.10	3.89	3.56	4.71	6.53	5.52	4.74

6 Conclusion

Although the unconditional distribution of exchange rate returns can only provide a rather limited set of 'stylized facts',[25] our results indicate that some of them are useful as benchmarks for financial market models, e.g. agent based models.

The mean of the return series just provides information on the trend for the sample considered. Thus, the large dispersion for this moment found for subsamples and the bootstrap is not surprising. In contrast, the unconditional variance appears to be a robust feature both in the bootstrap and under time aggregation (correcting for the averaging effect). The skewness appears to be close to zero for all return series considered. The bootstrap showing a large variance for this moment, it should not be considered as a benchmark. The empirical kurtosis appears to be a much more robust statistic with the exception of the first few years of our sample. Furthermore, under time aggregation, a decrease of excess kurtosis is found as reported in the literature. However, the Hill-estimator of the tail index indicates that the theoretical fourth moment might not be existing for daily and weekly exchange rate returns. Consequently, the use of the kurtosis and the Jarque-Bera statistic as a benchmark requires additional attention. Given that the Hill-estimator also provides robust estimates of the tail index with increasing values under time aggregation, this might be a more appropriate benchmark for the fat-taildness of exchange rate returns. Finally, the Kolmogorov-Smirnov statistic provides a robust non-parametric estimate of the deviation of the unconditional distribution from a normal distribution. Obviously, its use is not constrained to a comparison with a fixed parametric distribution, but it can also be used to assess differences between empirical distributions and distributions generated by the simulation of a model, e.g. an agent based model.

In addition to the analysis of statistics of the unconditional return distributions, the bootstrap also provides estimates of their standard deviations. This information can also be used when it comes to compare the empirical distribution with model based distributions when, e.g., Monte Carlo simulation results both in estimates of the statistics and their distribution.

Summarizing our findings it might be stated that from the set of statistics

[25]The more interesting aspects related to conditional distributions and long range dependence are discussed in Winker and Jeleskovic (2006, forthcomming).

on unconditional distributions considered in this paper, the (unconditional) variance, the Kolmogorov-Smirnov statistic and the tail index appear to be most promising as benchmarks for financial market models. Although the estimates of the empirical kurtosis appear to be quite robust and exhibit an interesting feature under time aggregation, its use required additional attention due to the possible non-existence of the theoretical fourth moment.

References

Alfarano, S., F. Wagner and T. Lux (2005). Estimation of agent-based models: The case of an asymmetric herding model. *Computational Economics* **26**(1), 19–49.

Arthur, W.B., B. LeBaron and R. Palmer (1999). Time series properties of an artificial stock market. *Journal of Economic Behavior & Organization* **23**, 1487–1516.

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

de Haan, L. and L. Peng (1998). Comparison of tail index estimators. *Statsitica Neerlandica* **52**, 60–70.

de Vries, C. C. (1994). Stylized facts of nominal exchange rate returns. In: *The Handbook of International Macroeconomics* (F. V. der Ploegg, Ed.). Blackwell. Cambridge, MA. pp. 348–389.

Drees, H. (1998). Optimal rates of convergence for estimates of the extreme value index. *Annals of Statistics* **26**, 434–448.

Embrechts, P. (2000a). Extreme value theory: Potentials and limitations as an integrated risk management tool. *Derivatives Use, Trading and Regulation* **6**, 449–456.

Embrechts, P., C. Klüppelberg and T. Mikosch (1999). *Modelling Extremal Events for Insurance and Finance.* Applications of Mathematics. 2nd ed. Springer. New York.

Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* **25**, 383–417.

Fama, E. (1991). Efficient capital markets: II. *Journal of Finance* **46**(5), 1575–1617.

Gençay, R. and F. Selçuk (2004). Extreme value theory and value-at-risk: Relative performance in emerging markets. *International Journal of Forecasting* **20**, 287–303.

Gençay, R., F. Selçuk and A. Ulugülyağci (2003a). High volatility, thick tails and extreme value theory in value–at–risk estimation. *Insurance: Mathematics and Economics* **33**, 337–356.

Gençay, R., F. Selçuk and A. Ulugülyağci (2003b). EVIM: a software package for extreme value analysis in Matlab. *Studies in Nonlinear Dynamics and Econometrics* **5**, 213–239.

Gilli, M. and E. Këllezi (2006). An application of extreme value theory for measuring financial risk. *Computational Economics* **27**(2–3), 207–228.

Gilli, M. and P. Winker (2003). A global optimization heuristic for estimating agent based models. *Computational Statistics and Data Analysis* **42**(3), 299–312.

Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics* **3**, 1163–1174.

Hsing, T. (1991). On tail estimation using dependent data. *Annals of Statistics* **19**(3), 1547–1569.

Jansen, D. W. and C. G. de Vries (1991). On the frequency of large stock returns: Putting booms and busts into perspective. *The Review of Economics and Statistics* **73**(1), 18–24.

Jenkinson, A. F. (1955). The frequency distribution of the annual maximum (minimum) values of meterological events. *Quarterly Journal of the Royal Meteorological Society* **81**, 158–172.

Kearns, P. and A. Pagan (1997). Estimating the density tail index for financial time series. *The Review of Economics and Statistics* **79**(2), 171–175.

Krämer, W. (2002). Statistische Besonderheiten von Finanzzeitreihen. *Jahrbücher für Nationalökonomie und Statistik* **222**(2), 210–229.

Lévy, P. (1924). Théorie des erreurs. La loi de Gauß et les lois exceptionnelles. *Bulletin de la SMF* **52**, 49–85.

Lilliefors, H. W. (1967). On the Kolmogorov–Smirnov test for normality with mean and variance unknown. *Journal of the American Statistical Association* **62**(318), 399–402.

Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *Journal of Economic Behavior & Organization* **33**, 143–165.

Lux, T. and M. Marchesi (2000). Volatility clustering in financial markets: a microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance* **3**, 675–702.

Lux, T. and S. Schornstein (2005). Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics* **41**, 169–196.

McNeil, A. J. (1997). Estimating the tails of loss severity distributions using extreme value theory. *ASTIN Bulletin* **27**, 1117–1137.

McNeil, A.J. and R. Frey (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* **7**(3–4), 271–300.

Quintos, C., Z. Fan and P. C. B. Phillips (2001). Structural change tests in tail behaviour and the Asian crisis. *The Review of Economic Studies* **68**(3), 633–663.

Resnick, S. and C. Starica (1998). Tail index estimation for dependent data. *Annals of Applied Probability* **8**, 1156–1183.

Segers, J. (2001). On the normal uniform local domain of attraction for intermediate order statistics. *Statistics and Probability Letters* **53**, 409–413.

Shleifer, A. (2000). *Inefficient Markets*. Oxford University Press. Oxford.

Tay, N.S.P. and S.C. Linn (2001). Fuzzy inductive reasoning, expectation formation and the behaviour of security prices. *Journal of Economic Behavior & Organization* **25**, 321–361.

von Mises, R. (1954). La distribution de la plus grande de n valeurs. In: *Selected papers of Richard von Mises* (P. Frank, Ed.). pp. 271–294. American Mathematical Society. Providence, RI. Reprint (first published in Rev. Math. Union. Interbalcanique, **1**, 141–160, 1936).

Wagner, N. and T. Marsh (2003). Measuring tail thickness under GARCH and an application to extremal exchange rate changes. Technical report. Haas School of Business, U.C. Berkeley.

Winker, P. and V. Jeleskovic (2006, forthcomming). Conditional distribution of – and long memory in – exchange rate returns: Statistics, robustness, time aggregation. Technical report. University of Giessen.

A Further Results

Figure 7: Hill plots for exchange rate times series

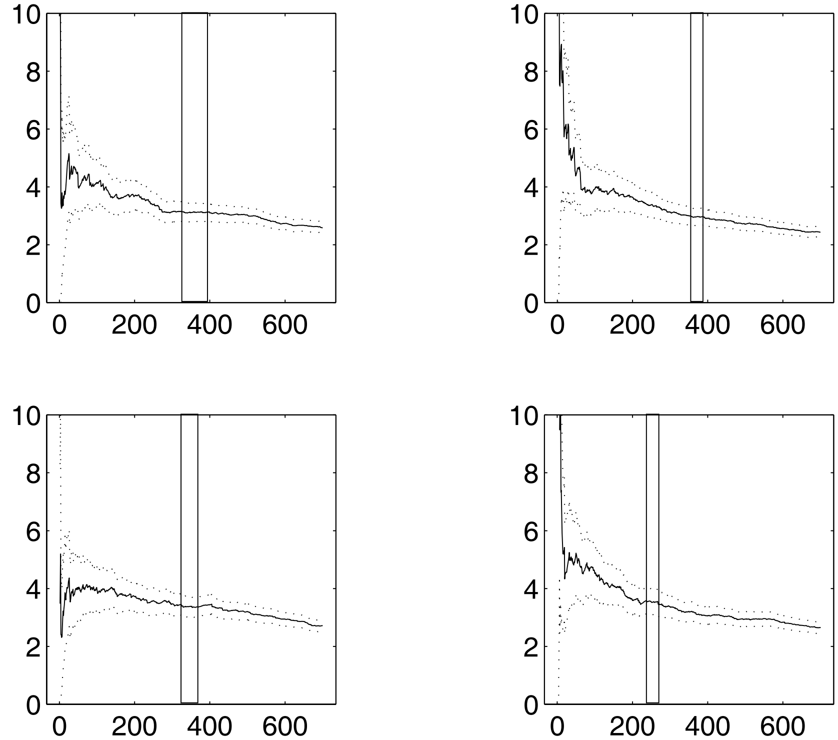

B Further Bootstrap Results

Table 9: Bootstrap Distribution of r_{FRF}

Statistic	FRF/US	Mean	5%	50%	95%	Std.dev.
Mean	0.11E-4	0.08E-4	-1.16E-4	0.06E-4	1.25E-4	0.74E-4
Std.dev.	0.0064	0.0064	0.0062	0.0064	0.0066	0.0001
Skewness	0.0999	0.0918	-0.1334	0.0762	0.3336	0.1494
Kurtosis	7.1693	7.0970	5.5974	6.9919	8.8993	1.0034
JB-Statistic	5479	5633	2128	5029	11054	2808
KS-Statistic	0.0544	0.0565	0.0511	0.0565	0.0624	0.0034
Tail index (left)	3.4041	3.4228	3.1540	3.4206	3.7015	0.1642
Tail index (right)	3.2916	3.3132	3.0531	3.3043	3.5880	0.1670
BMM α_{left}	1.5607	-6.3204	-42.7591	-4.8036	31.0540	86.6562
BMM α_{right}	1.5750	1.1486	-4.5699	3.2175	21.8487	83.2922
POT α_{left}	5.4496	2.9790	-36.4431	-2.2808	35.1590	158.8977
POT α_{right}	3.0335	2.9195	-9.1539	2.7098	20.0833	49.0743

Table 10: Bootstrap Distribution of r_{GBP}

Statistic	GBP/US	Mean	5%	50%	95%	Std.dev.
Mean	0.26E-4	0.23E-4	-0.93E-4	0.26E-4	1.33E-4	0.70E-4
Std.dev.	0.0061	0.0061	0.0060	0.0061	0.0062	0.0001
Skewness	0.1300	0.1275	-0.0814	0.1310	0.3165	0.1177
Kurtosis	6.6767	6.6636	5.9048	6.6301	7.5103	0.5044
JB-Statistic	4272	4338	2692	4165	6452	1183
KS-Statistic	0.0633	0.0647	0.0593	0.0645	0.0707	0.0035
Tail index (left)	3.6054	3.6166	3.3100	3.6137	3.9429	0.1989
Tail index (right)	3.2611	3.2740	3.0249	3.2703	3.5195	0.1535
BMM α_{left}	1.5650	-18.8731	-41.0539	2.6932	39.5610	657.3083
BMM α_{right}	1.5608	-6.3126	-32.3930	-4.4989	12.8303	45.8525
POT α_{left}	3.7721	3.6553	-28.2655	1.8700	33.9194	205.2695
POT α_{right}	9.6023	2.3739	-36.5203	-3.2861	49.3620	136.5969

Table 11: Bootstrap Distribution of r_{CHF}

Statistic	CHF/US	Mean	5%	50%	95%	Std.dev.
Mean	-1.08E-4	-1.09E-4	-2.52E-4	-1.09E-4	0.37E-4	0.87E-4
Std.dev.	0.0076	0.0075	0.0073	0.0075	0.0078	0.0002
Skewness	0.0438	0.0611	-0.6272	0.0298	0.8307	0.4374
Kurtosis	13.6992	13.3188	5.7076	13.3847	22.1348	4.8311
JB-Statistic	36009	41071	2324	33963	11521	35901
KS-Statistic	0.0527	0.0540	0.0471	0.0538	0.0608	0.0042
Tail index (left)	3.4652	3.4790	3.2043	3.4784	3.7515	0.1629
Tail index (right)	3.5035	3.5230	3.2119	3.5139	3.8462	0.1917
BMM α_{left}	1.7265	-0.7792	-6.9730	1.5959	8.2941	45.4457
BMM α_{right}	3.1652	3.0660	1.5629	1.5980	10.7893	17.0052
POT α_{left}	3.3183	2.6697	-14.7760	2.1455	9.2199	75.1262
POT α_{right}	2.1223	3.1704	-3.8357	1.9129	12.2847	38.2573

Table 12: Bootstrap Distribution of r_{JPY}

Statistic	JPY/US	Mean	5%	50%	95%	Std.dev.
Mean	-1.42E-4	-1.44E-4	-2.75E-4	-1.47E-4	-0.20E-4	0.78E-4
Std.dev.	0.0066	0.0066	0.0064	0.0066	0.0068	0.0001
Skewness	-0.4862	-0.4889	-0.7094	-0.4804	-0.2998	0.1277
Kurtosis	7.3763	7.3464	6.2287	7.2704	8.7019	0.7542
JB-Statistic	6320	6441	3431	6035	10818	2296
KS-Statistic	0.0705	0.0719	0.0654	0.0717	0.0787	0.0040
Tail index (left)	3.0992	3.1114	2.8427	3.1110	3.3694	0.1625
Tail index (right)	3.3893	3.4113	3.1485	3.4099	3.6669	0.1620
BMM α_{left}	1.5720	-6.4014	-35.0290	1.5927	48.9208	258.5851
BMM α_{right}	1.5608	-0.0056	-25.6277	-0.8545	27.9452	40.7210
POT α_{left}	3.0420	3.1576	-50.5164	2.0072	44.3912	181.7041
POT α_{right}	7.5523	3.1911	-7.4871	-0.9743	-0.0359	104.2892

Figure 8: Bootstrap Distributions (Block Length 20 Days)