The tail fatness of FX returns reconsidered

Citation for published version (APA):

Palm, F. C., Huisman, R., Koedijk, C. G., & Kool, C. J. M. (2002). The tail fatness of FX returns reconsidered. *De Economist*, 150(3), 299-312. <https://doi.org/10.1023/A:1016165808823>

Document status and date:

Published: 01/01/2002

DOI:

[10.1023/A:1016165808823](https://doi.org/10.1023/A:1016165808823)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

NOTES AND COMMUNICATIONS

THE TAIL-FATNESS OF FX RETURNS RECONSIDERED

1 INTRODUCTION

It is a well-known stylized fact that foreign exchange returns are non-normal and tend to have fat-tailed distributions. Although information about the precise magnitude of the tail-fatness is crucial for applications such as risk analysis, little consensus exists in this respect, due to estimation problems. Typically, empirical estimation of tail-fatness is conditional on the specific (fat-tailed) distribution – such as the Student-t or stable – chosen.¹ Since these alternative distributions are non-nested, an incorrect choice of distribution of returns may lead to significant estimation errors and flawed inference.

In recent years, extreme value analysis has been proposed to overcome this potential difficulty. It investigates the distribution of the maximum (minimum) in large samples, thereby determining the shape of the tails of a distribution. The limit law for the maximum is characterized by the so-called tail-index α , which corresponds one-to-one with the number of existing moments of the underlying distribution. Note that alternative distributions like the stable and Student-t now are nested within the limit law for extremes; for the stable distribution, the characteristic exponent equals the tail-index α (< 2), while for the Student-t distribution it is equal to the number of degrees of freedom ν . The gain in applying extreme value analysis is that one can nest and test for different tail sizes. However, information about the center characteristics of the distribution is disregarded. Since it is information about the outlier observations that is crucial for studying issues like exchange rate volatility, exchange rate risk and value at risk, the use of extreme value analysis offers a positive net gain in our view.

The best-known and most often applied extreme value estimator for the tail-index was proposed by Hill (1975). Although asymptotically unbiased, this maximum likelihood estimator suffers severely from small sample bias. In fact, this holds for most alternative estimators, see Pictet et al. (1996). Consequently, the

¹ See Westerfield (1977), Rogalski and Vinso (1978), Boothe and Glassman (1987), and McCulloch (1997).

empirical applicability of tail-index estimation is limited to cases of extremely large samples, either in the form of high frequency data or a very long sample period.² In real life, often neither of these two conditions is fulfilled and only a small sample analysis is feasible.

In this contribution, we re-examine the tail-fatness of exchange rate returns using the tail-index estimator proposed by Huisman, Koedijk, Kool, and Palm (2001) – henceforth HKKP –, which is unbiased in small samples. We show that previous tail estimates in the literature overestimate the amount of tail-fatness in exchange rate returns. Additionally, goodness-of-fit statistics provide supportive evidence for using the Student-t distribution as an appropriate (local) approximation of the tail probability of exchange rate returns. The evidence is more convincing for floating than for exchange rates which have been subject to the exchange rate mechanism (ERM) within the European Monetary System (EMS).

The paper is set up as follows. In section 2, we briefly discuss extreme value theory and introduce the HKKP estimator to obtain unbiased tail-index estimates in small samples. In section 3, we describe data on sixteen exchange rates between January 1979 and April 1996. In section 4, we apply the HKKP tail index estimator and explicitly compare the HKKP tail index estimates with results from the existing literature. Also, we successfully examine the goodness-of-fit of a Student-t distribution with the degree of freedom equalling the HKKP tail estimates. Conclusions are given in section 5.

2 THE TAIL-INDEX

The tail-index α measures the speed with which the tail of a fat-tailed distribution $F(\cdot)$ approaches zero, when $F(\cdot)$ fulfills the following regular variation condition at infinity³:

$$\lim_{t \rightarrow \infty} [1 - F(tx)] / [1 - F(t)] = x^{-\alpha}, \quad (1)$$

which implies that the higher α , the less fat-tailed the distribution is. Because the tail-index α equals the maximum number of finite moments in the sample, the tail-index can directly be used to test for the number of existing moments.

The HKKP estimator is a simple and appropriate tool to obtain unbiased tail-index estimates in small samples. The methodology starts from the observation that the bias of the conventional Hill (1975) estimator is a function of sample

2 Recently, efficient non-parametric estimators have been proposed, see Beirlant et al. (1996), for example. However, their implementation is complex and small-sample properties are as yet unknown.

3 Many studies have shown that many financial series, including exchange rate returns fulfill this condition.

size n . Then, the bias function is used to correct for the small sample bias.⁴ Suppose a sample of n observations is drawn from some unknown fat-tailed distribution. Let the parameter γ be the tail-index⁵ of this distribution, while x_i is the i th increasing order statistic ($i = 1, 2, \dots, n$). Hill (1975) proposes the following estimator for the tail-index:

$$\gamma(k) = \left[\sum_{j=1}^k \ln(x_{n-j-1}) \right] / k - \ln(x_{n-k}), \quad (2)$$

where k is the pre-specified number of tail observations to include ($k = 1, \dots, n-1$). The choice of k is crucial to obtain an unbiased estimate of the tail-index.⁶

HKKP (2001) show that for a quite general class of distribution functions, the bias of the Hill estimator increases monotonically with k , while the variance is proportional to $1/k$. A trade-off results between unbiasedness and accuracy. Most empirical studies suffer from this trade-off problem. Generally, a single ‘optimal’ k is selected. However a bias exists for any k exceeding zero. The dilemma is resolved when the sample size goes to infinity for given k as then the bias goes to zero.

The HKKP estimator overcomes the problem of the need to select an optimal k in small samples, by exploiting an important characteristic of the bias function. For values of k smaller than some threshold value κ (for example $k \leq n/2$), the bias of the conventional Hill estimate of γ increases almost linearly in k and can be approximated by:

$$\gamma(k) = \gamma + \beta k + \varepsilon(k), \quad k = 1, 2, \dots, \kappa, \quad (3)$$

where $\varepsilon(k)$ is a disturbance term. Instead of selecting a single value of k to estimate the tail-index of the distribution under consideration, we propose to compute $\gamma(k)$ for a range of values of k from 1 to κ . Subsequently, the vector of the computed $\gamma(k)$ ’s is used in equation (3). We already argued that an unbiased estimate of γ can be obtained only for $k \rightarrow 0$. Evaluation of equation (3) for k approaching zero yields the bias-corrected estimator of γ proposed by HKKP (2001).

Although the coefficients in (3) can be estimated by Ordinary Least Squares (OLS), two issues have to be considered. First, as mentioned, the variance of Hill

4 This estimator, which may be seen as a modified Hill estimator, is studied in more detail in Huisman et al. (2001).

5 By definition, γ equals $1/\alpha$, where α refers to the maximum number of existing finite moments. In the literature, the tail-index is referred to as either α or γ ; we use both interchangeably.

6 Various studies select k to minimize the MSE of the Hill estimates conditional on a known distribution, see for instance Koedijk and Kool (1994). See Loretan and Phillips (1994) for an alternative approach.

estimates $\gamma(k)$ varies with $1/k$, so that the error term $\varepsilon(k)$ in equation (3) is heteroskedastic. As an alternative, we suggest a Weighted Least Squares (WLS) approach. Multiplying the right- and left-hand sides of equation (3) by $\sqrt{k}$ and applying OLS to the transformed equation yields the WLS estimator which accounts for heteroskedasticity. The estimate of γ from the WLS regression is an approximately unbiased estimate of the tail-index γ .

Second, the disturbances in (3) are correlated due to the construction of $\gamma(k)$. The variables $\gamma(k)$, say for $k = s$ and $k = m$, are correlated, since they are based on order statistics, which are correlated, and they have $1 + \min(s, m)$ common observations, see equation (2). Consequently, the usual formulae for the standard errors are inappropriate for the OLS and WLS estimators, but the well-known formulae (see HKKP (2001) for more details) for the standard errors of these estimators when applied to the generalized regression model are consistent and will be used in the sequel. In the empirical part, we report weighted least squares estimates⁷ and appropriate standard errors accounting for the disturbance correlation resulting from the use of order statistics and overlapping observations as mentioned above.

3 DATA

The data consist of London closing mid-prices against the pound sterling for the period January 2, 1979 and April 4, 1996 obtained from DATASTREAM. The currencies of sixteen countries are considered: Austria, Belgium, Canada, Denmark, France, Germany, Ireland, Italy, Japan, the Netherlands, Norway, Spain, Sweden, Switzerland, the United Kingdom and the United States. We sample the data at the weekly frequency⁸ (using end-of-period data) to reduce the influences of ARCH effects and other types of dependency between subsequent observations that are present in many financial data, since the tail-index estimates are produced assuming independent observations⁹. We have preferred not to use more recent data as the probability distribution of EMS exchange rates in more recent years is likely to have been affected by the expected start of the common currency euro as of January 1, 1999 and the decisions taken in the years preceding

7 HKKP (2001) have found the WLS estimator to be more accurate than the OLS estimator and to be hardly less accurate than the generalized least squares estimator which also accounts for the correlation between elements of the disturbance term in (3).

8 Kearns and Pagan (1997) argue that the dependency between the observations results in serious biases for GARCH and IGARCH type models, especially when those estimates are obtained from small samples. Lucas (1997) shows that the tail-index estimator that we use in this study produces rather unbiased tail-index estimates for a GARCH(1,1) process in a sample of 250 observations. This, combined with the fact that we use weekly data to eliminate as much dependency as possible, supports that the results presented here are accurate.

9 Note that the true tail-index of the underlying return distribution is unaffected by time-aggregation; thus, the choice of sampling frequency should only affect the estimation results through the change in number of observations.

TABLE 1 – SUMMARY STATISTICS

Currency	numeraire: German Mark				numeraire: US dollar			
	mean	std.	skew	kurt.	Mean	std.	skew	kurt.
Austria	−0.01	1.17	−0.12	3.55	−0.63	5.22	−0.10	1.47
Belgium	0.35	1.54	1.66	20.72	−0.28	5.17	−0.09	1.52
Canada	0.96	5.45	0.02	1.50	0.33	2.18	0.09	2.67
Denmark	0.56	1.67	1.53	11.76	−0.07	5.09	−0.10	1.15
France	0.78	1.85	2.86	28.21	0.15	5.14	0.06	1.86
Germany	–	–	–	–	−0.63	5.25	−0.07	1.46
Ireland	0.93	2.52	1.76	16.86	0.30	5.17	0.21	2.66
Italy	1.57	3.25	2.28	18.09	0.95	5.04	0.68	5.86
Japan	−0.54	4.47	−0.25	1.31	−1.17	4.99	−0.54	2.77
Netherlands	0.10	0.87	0.26	13.38	−0.53	5.17	−0.09	1.27
Norway	0.80	2.67	2.09	15.33	0.17	4.48	0.45	3.22
Spain	1.47	3.92	8.69	148.89	0.84	5.57	2.96	33.49
Sweden	1.11	3.94	6.32	84.03	0.48	4.95	3.15	34.10
Switzerland	−0.23	2.50	0.15	3.11	−0.86	6.00	−0.12	1.30
U. Kingdom	0.96	3.96	0.59	3.03	0.33	5.16	0.24	3.38
U. States	0.63	5.25	0.07	3.46	–	–	–	–

The mean, standard deviation, skewness and excess kurtosis of the weekly returns in annualized percentages are provided for the period from January 2, 1979 through April 4, 1996.

this date to make that start possible. All data are expressed as the number of foreign currencies to pay for one unit of the numeraire currency.

Table 1 contains summary statistics of the weekly data. With respect to skewness and kurtosis the following observations can be made from the table. First, the EMS currencies denoted in German Mark (DM) are skewed to the right, whereas the amount of skewness is much smaller and hovers around zero for the floating exchange rates denoted in US dollars. Second, all exchange rate returns exhibit excess kurtosis,¹⁰ reflecting apparent fat-tailedness. The excess kurtosis is most apparent for the EMS exchange rates in general, and for the Spanish peseta and Swedish krone in particular.

4 RESULTS

In this section, we present tail-index estimates α for the 16 currencies described in the previous section with the US dollar and the DM alternating as benchmark currency. We compare the tail-index estimates obtained from the HKKP estimator

10 Excess with respect to the normal distribution which has a kurtosis equal to 3.

with results from the literature that are mostly obtained using maximum likelihood techniques. Subsequently, goodness-of-fit tests are used to investigate whether a hypothesized Student-t distribution with degrees of freedom equal to the obtained HKKP tail-index estimates adequately reproduces the actual distribution of exchange rate returns.

4.1 Tail-Index Estimates

The observed excess kurtosis in Table 1 already supports the findings of many other studies that the distribution of foreign exchange rate returns is fat-tailed. Note though that the kurtosis estimate is conditional on existence of the fourth moment of the underlying distribution. Especially for cases where quite large excess kurtosis values are obtained, such as, for example, the Spanish peseta and the Swedish krone against the DM, this assumption might be unwarranted. A direct way to test for the existence of the fourth moments is to evaluate the tail-index estimates for each exchange rate. For this purpose, HKKP tail estimates with appropriate standard errors are reported in Table 2. Estimates are shown for left-tail and right-tail observations separately and for the combination of both tails.¹¹

According to Table 2, the fourth moments are likely to exist for all exchange rates *vis-à-vis* the US dollar. Considering both tails simultaneously, tail-index estimates are found to vary between 4.0 for the Spanish peseta and 8.1 for the British pound. The tail-index estimates for currencies *vis-à-vis* the German Mark are smaller, indicating a higher amount of tail-fatness. The estimates vary between 2.2 for the French franc and 5.4 for the Canadian dollar. For many of these exchange rates, the standard errors are such that the finiteness of the fourth moments can be rejected.

From Table 2 we also conclude that the tail-index of the left and the right tail often differ considerably. For example, we obtain a tail-index estimate equal to 4.3 for the left tail of Canadian *vis-à-vis* the German Mark and an estimate equal to 7.3 for the right tail. If these tails differ significantly, one should use tail-specific estimates when drawing inference from the tails. The columns headed 't' contain the t-statistics for testing the hypothesis that the left tail-index equals the right tail-index (accounting appropriately for the dependence between left and right tail-index estimates as they are functions of order statistics; see HKKP (2001) for details). The results show that the difference between the left tail-index and the right tail-index is insignificant for the majority of exchange rates. It is significant at the 5% level for the Canadian dollar, the Danish krone, the Norwegian krone, the Spanish peseta and the Swedish krone all *vis-à-vis* the German

¹¹ All observations are taken in deviation of their sample mean. The left tail is examined using the absolute value of all negative returns, the right tail is examined using all positive returns and to examine both tails simultaneously we use the absolute values of all returns.

TABLE 2 – TAIL-INDEX α ESTIMATES

Currency	numeraire: German Mark				numeraire: US Dollar			
	Both	Left	Right	t	Both	Left	Right	t
Austria	3.42 (0.33)	3.46 (0.48)	3.19 (0.45)	0.42	5.75 (0.56)	5.62 (0.80)	5.99 (0.82)	−0.32
Belgium	3.25 (0.32)	3.95 (0.54)	2.72 (0.39)	1.86	5.38 (0.52)	5.12 (0.72)	5.70 (0.79)	−0.54
Canada	5.45 (0.53)	4.37 (0.61)	7.26 (1.02)	−2.43	4.61 (0.45)	4.01 (0.55)	5.66 (0.80)	−1.70
Denmark	3.29 (0.32)	4.16 (0.56)	2.70 (0.39)	2.14	5.28 (0.52)	4.84 (0.68)	5.72 (0.80)	−0.84
France	2.21 (0.22)	2.71 (0.34)	2.49 (0.39)	0.42	5.24 (0.51)	5.52 (0.76)	4.77 (0.67)	0.74
Germany	–	–	–	–	5.32 (0.52)	5.94 (0.83)	4.83 (0.67)	1.04
Ireland	2.62 (0.26)	2.89 (0.38)	2.76 (0.42)	0.23	5.76 (0.56)	5.44 (0.74)	6.01 (0.86)	−0.50
Italy	2.28 (0.22)	2.36 (0.30)	2.43 (0.37)	−0.13	4.71 (0.46)	5.28 (0.72)	4.23 (0.61)	1.12
Japan	5.25 (0.51)	5.55 (0.78)	5.30 (0.73)	0.22	6.10 (0.59)	9.04 (1.32)	5.80 (0.78)	2.12
Netherlands	2.79 (0.27)	2.71 (0.38)	2.86 (0.40)	−0.27	5.37 (0.52)	6.28 (0.88)	4.71 (0.65)	1.44
Norway	3.31 (0.32)	4.95 (0.66)	2.51 (0.37)	3.21	4.44 (0.43)	4.89 (0.67)	4.04 (0.57)	0.97
Spain	3.19 (0.31)	4.60 (0.61)	2.53 (0.37)	2.90	4.03 (0.39)	5.08 (0.69)	3.28 (0.47)	2.16
Sweden	3.03 (0.30)	4.07 (0.55)	2.43 (0.36)	2.52	4.26 (0.42)	5.08 (0.70)	3.59 (0.51)	1.72
Switzerland	4.79 (0.47)	5.61 (0.79)	4.18 (0.58)	1.46	5.73 (0.60)	8.31 (1.16)	4.55 (0.64)	2.85
U. Kingdom	3.85 (0.38)	4.16 (0.56)	4.02 (0.59)	0.18	8.11 (0.79)	8.03 (1.09)	7.79 (1.12)	0.15
U. States	5.23 (0.52)	4.83 (0.67)	5.94 (0.83)	−1.04	–	–	–	–

Tail-index estimates for the weekly returns for the period from January 2, 1979 through April 4, 1996 are given for both tails jointly and for the left and the right tails separately. The standard errors given in parentheses account for the correlation resulting from the use of order statistics and overlapping observations. The two-sided t-test of the equality of the tails for the exchange rate returns in terms of DM and US dollar respectively is given in the column 't'.

Mark and for the Japanese yen, the Spanish peseta, and the Swiss franc *vis-à-vis* the US dollar.

4.2 The Tail-Index of EMS Rates Versus Floating Rates

Koedijk et al. (1992) and Koedijk and Kool (1994) show that EMS exchange rate returns exhibit fatter tails than floating rate returns. The information in Table 2 provides clear evidence in support of this conclusion. For all EMS currencies, tail estimates are lower when measured *vis-à-vis* the DM than *vis-à-vis* the US dollar, regardless of whether the left-tail, right-tail or both are considered. Table 3 contains the results of formally testing the null hypothesis that the tail-index of a currency *vis-à-vis* the German Mark equals the tail-index of the same currency but now against the US dollar. The test statistic is asymptotically normally distributed. For negative values the tail fatness against the DM exceeds that against the US dollar.

For most of the countries that have floated against both the US dollar and the DM, the equality of tail-fatness across numeraire currencies cannot be rejected. This is the case for Canada, Japan, Spain, and Switzerland. Note that Spain has participated in the ERM only from 1989 onward, and has devalued a few times

TABLE 3 – EMS VERSUS FLOATING

Currency	Both	Left	Right
Austria	-3.58	-2.33	-2.99
Belgium	-3.49	-1.30	-3.39
Canada	1.22	0.45	1.23
Denmark	-3.28	-0.78	-3.42
France	-5.48	-3.37	-2.94
Germany	-	-	-
Ireland	-5.10	-3.07	-3.40
Italy	-4.77	-3.74	-2.54
Japan	-1.08	-2.28	-0.47
Netherlands	-4.37	-3.73	-2.41
Norway	-2.09	0.07	-2.25
Spain	-1.69	-0.53	-1.26
Sweden	-2.41	-1.14	-1.86
Switzerland	-1.25	-1.93	-0.43
U. Kingdom	-4.88	-3.16	-2.98
U. States	-	-	-

We report t-statistics for the two-sided test of the equality of the tails for exchange rate returns in terms of DM and US dollar respectively. 'Both' refers to the test on both tails simultaneously. 'Left' and 'Right' refer to the left and the right tail respectively.

against the DM in 1992 and 1993. For all ERM countries equality of tail-indices across numeraire currencies is firmly rejected with the tails against the DM showing significantly more fatness than the tail-indices against the US dollar. The same holds (marginally) for Norway and Sweden, even though they did not formally participate in the EMS over the sample period. Consequently, the HKKP estimates support earlier conclusions that the return distribution of non-floating exchange rates has substantially fatter tails than the distribution under floating.

4.3 *A Comparison: HKKP Estimates Versus the Literature*

The HKKP tail-index estimates reported in Table 2 are much higher than values found in previous studies for floating rates. For example, Loretan and Phillips (1994) report α 's for the DM / US dollar rate around 3.4 for the period 1978-1991, while we find 5.3 for both tails instead. Boothe and Glassman (1987) report a tail-index estimate of 3.2 for the German Mark / US dollar rate for the period 1973-1984.

These differences might be due to the different sample periods. However for roughly the same sample period as used by Loretan and Phillips, HKKP (2001) find systematically higher tail-index estimates α for all currencies examined by Loretan and Phillips (1994).

An alternative explanation focuses on the shortcomings of previously applied procedures. Loretan and Phillips (1994) obtain their estimates from the conventional Hill estimator (2) which suffers from small sample bias. Boothe and Glassman (1987) use a maximum likelihood procedure to estimate the tail index conditional on the true distribution being Student-t. As argued before, this method suffers from bias as well in small samples and the distributional assumption may be incorrect and lead to estimation errors. The inequality of left and right tails for many currencies, as illustrated in Table 2, is of additional concern. Assuming a symmetric distribution like a Student-t or a sum-stable distribution definitely leads to misspecification. Overall, we conclude the HKKP results must be preferred. Consequently, exchange rate returns have fatter tails than the normal distribution, but they exhibit considerably less tail-fatness than has been argued in the literature before. This holds especially true for floating exchange rate returns.

4.4 *Distributional Characteristics*

The fact that the tail-fatness of exchange rate returns has been considerably overestimated in the literature has important implications for the issue of which statistical distribution captures the observed distribution of exchange rate returns best. Our results, for instance, indicate that the stable distribution with $2 \geq \alpha$ is likely to be inappropriate for exchange rate returns. The characteristic exponent of the stable distribution which equals the tail-index is by definition below 2, while we find tail-index values around 5 for most floating exchange rate returns.

Westerfield (1977), Rogalski and Vinso (1978), and Boothe and Glassman (1987) examined several distributions to capture tail-fatness. The Student-t distribution, a mixture of normals, and the sum-stable distribution are shown to provide the best fit among various alternative candidates, although goodness-of-fit tests only provide weak statistical evidence. This may be due to the poor quality of previous tail estimates.

Here, we investigate whether a (possibly asymmetric) Student-t distribution with number of degrees of freedom equal to the estimated tail-indices, adequately fits the observed frequencies in the tails of the exchange rate return distribution.

Let x_i be the exchange rate returns under consideration. To fit the Student-t distribution, we assume that $y_i = \theta(x_i - m)$ is Student-t distributed with α degrees of freedom. Here m is a location parameter that equals the sample mean. The scale parameter θ is a function of the variance and can be derived as follows. For a Student-t distribution with α degrees of freedom, the variance of y_i equals $\alpha/(\alpha-2)$ where $\alpha > 2$. Consequently, the scale parameter θ can be expressed in terms of the variance of the exchange rate and the number of degrees of freedom α as: $\theta = \sqrt{[\alpha/(\alpha-2)\text{var}(x_i)]}$. Next, we set the number of degrees of freedom equal to the HKKP tail-index estimate for each currency. As the tail-index estimates for the left and right tails differ significantly for some exchange rates, the whole distribution is fitted by separately using the estimated tail-index for each tail. The above procedure allows for the estimation of the distribution for exchange rate returns conditional on the HKKP tail estimates and the sample mean and variance.

To test the hypothesis that the Student-t distribution is a good approximation for the unconditional distribution of exchange rate returns, we apply a Pearson goodness-of-fit test similar to that used by Boothe and Glassman (1987). The goodness-of-fit test compares the observed and expected number of observations in c intervals.¹² We set the number of intervals c equal to 20.¹³ The intervals (except the first in the left tail and the last in the right tail which are open to the left and the right respectively) are chosen such that they are of equal length and that the tail intervals have 5 expected observations, which is the minimum for the test statistic to be chi-squared distributed.

Table 4 contains the results both for the whole distribution and for the tails. With respect to the latter, we distinguish between fitting the 5% largest and 5% smallest observations and fitting the 10% largest and 10% smallest observations.

¹² When the theoretical distribution is given under the null hypothesis, the test statistic is asymptotically chi-square distributed with $(c-1)$ degrees of freedom. When the number of estimated parameters is k , the statistic is bounded between chi-square distributions with $(c-1)$ and $(c-k-1)$ degrees of freedom, respectively.

¹³ As in Boothe and Glassman (1987), who use 14 intervals, our results are robust with respect to the number of intervals chosen. See Bain and Engelhardt (1987) for details on goodness-of-fit tests.

TABLE 4 – GOODNESS-OF-FIT STATISTICS FOR STUDENT-T

Currency	numeraire: DM				numeraire: US dollar		
	5%	5% (R)	10%	all	5%	10%	All
Austria	9.49	10.82	18.57	22.95	4.87	22.87	20.82
Belgium	18.13	19.39	31.06	44.00	5.49	13.46	13.88
Canada	12.66	14.96	24.29	23.27	6.36	12.06	18.20
Denmark	18.29	11.17	24.94	46.15	7.75	15.14	12.39
France	21.79	18.47	37.92	100.80	5.30	16.57	20.58
Germany	–	–	–	–	11.35	16.81	21.39
Ireland	17.89	16.01	25.64	77.83	4.89	15.03	16.17
Italy	36.18	44.27	48.64	92.67	9.81	7.75	14.87
Japan	10.59	10.38	13.17	18.84	13.29	20.65	45.33
Netherlands	12.26	15.80	26.58	28.46	7.57	16.32	12.75
Norway	46.62	23.64	59.06	80.01	10.28	19.00	26.83
Spain	46.51	47.80	51.88	118.31	10.57	12.29	48.07
Sweden	21.27	22.69	40.36	74.58	15.82	16.93	32.82
Switzerland	3.46	5.65	13.99	10.74	16.28	23.59	33.52
U. Kingdom	27.57	32.56	33.33	41.45	7.57	14.79	27.68
U. States	11.21	13.09	17.28	21.50	–	–	–

The goodness-of-fit statistics test the fit of the Student-t distribution with the number of degrees of freedom equal to the tail-index estimate α . The tests are performed on all observations (all), on the 10% largest and the 10% smallest observations (10%), on the 5% largest and the 5% smallest observations (5%) and on the 5% largest and smallest obtained from the sample from which data for all realignment periods are excluded 5% (R). The expected and observed frequencies are calculated over 20 intervals for all and 10%; they are calculated over 10 intervals in case of 5%. Separate tail estimates are used for the right and left tails. The chi-square critical value (neglecting the effect of parameter estimation) for 19 degrees of freedom at the 5% level is 30.14; for 9 degrees of freedom it equals 16.92

From Table 4, we observe that the Student-t is a good approximation for the complete distribution of returns for many of the floating exchange rates (see the ‘all’ column). For the tails, the fit is even better.¹⁴ For all currencies against the US dollar, the test fails to reject the Student-t distribution both for the 5% and 10% categories. Note that this result also holds for the exchange rates for which the Student-t distribution was rejected for all observations: the Japanese yen, the Spanish peseta, the Swedish krone and the Swiss franc. This result implies that one could rely on the Student-t distribution to obtain tail information for these

¹⁴ Especially the fit of the tails is relevant here. Although the Student-t fits the overall distribution rather well too, if the goal is to fit the complete distribution, empirical distribution kernel estimates and splines are likely to produce an even better fit.

floating rates in practice. Also, the Student-t distribution has the advantage of providing accurate exceedence probabilities or quantiles for the tails in a straightforward way. Especially for events that occur with a probability smaller than $1/n$, a good fit of the tails through a Student-t distribution is potentially valuable. Extrapolation out-of-sample using the functional form of the Student-t can be done with great ease and could be very useful for risk management.

The results are not that strong for the exchange rates *vis-à-vis* the German Mark. The Student-t is rejected to fit the tails for the Belgian franc, the Danish krone, the French franc, the Irish pound, the Italian lira, the Norwegian krone, the Spanish peseta, the Swedish krone, and the British pound. One reason why the Student-t distribution is rejected in the tails for the exchange rates mentioned above may be the realignments that took place in the EMS. To test for the impact of these realignments we computed HKKP tail-index estimates from a sample where all realignment observations were excluded. Fitting a Student-t with the new tail estimate for the 5% smallest and 5% largest observations leads to goodness-of-fit statistics in the 5% (R) column. The results improve somewhat and show that we cannot reject the Student-t distribution for the Danish krone and possibly the Irish pound. Note that the Student-t is rejected to fit the whole distribution but accepted to fit the 10% tail for the Danish krone, the Irish pound, and the Dutch guilder. Clearly, realignments do not provide the complete story about the difference between ERM and floating exchange rate return distributions.

5 CONCLUSIONS AND PRACTICAL IMPLICATIONS

In this contribution, we focus on the tail characteristics of the unconditional distribution of exchange rate returns. It is a stylized fact that exchange rate returns are fat-tailed. However, methods used to compute the degree of tail-fatness generally suffer from small sample bias. We apply the unbiased HKKP tail-index estimate, which essentially is a modified Hill estimate corrected for small sample bias. The correction exploits the fact that the bias is approximately linear in the number of tail observations used.

We apply this method to currencies of 16 countries, measured against both US dollar and DM for the period 1979-1996. We conclude first, that the magnitude of tail fatness in exchange rate returns has been seriously over-estimated in the past, especially for many floating rates. Second, the tail-indices for the left tail (negative returns) differ significantly from those of the right tail for many exchange rates; the assumption of an underlying symmetric distribution is inappropriate, therefore. Third, the tails of EMS exchange rate returns are significantly fatter than the tails of floating rate returns. Fourth, the Student-t distribution is found to fit the tails of all floating exchange rates when the degree of freedom parameter is set equal to the HKKP tail estimate. For some of the exchange rate returns in terms of the German Mark apparent skewness and realignments in the

EMS result in somewhat poorer fits. This conclusion holds even when realignments are excluded. More research is needed to better understand the difference between exchange rates that are subject to an exchange rate mechanism and floating rates in this respect. Nevertheless, the Student-t fits the tails well for almost all exchange rates under consideration.

The results put forward in this contribution have important implications for managing exchange rate risk. Risk on extreme events on floating exchange rate positions has been seriously over-estimated in the past. Furthermore, one needs to analyze positive and negative returns separately and one cannot directly assume that the probability of extreme events is the same for short and long foreign exchange positions. Also, the probability on extreme events is much higher for EMS rates.

Finally, the Student-t distribution with number of degrees of freedom set equal to the HKKP tail-index estimate can be used as an accurate approximation for tails of the unconditional return distribution of floating exchange rates and of most EMS exchange rates with the German Mark as numeraire. Exceedence probabilities and quantile estimates used as inputs for Value at Risk formulae can be easily obtained by extrapolation of the fitted Student-t distribution. This is especially attractive for the analysis of the risk on out-of-sample events which occur with a probability smaller than $1/n$.

Ronald Huisman*
Kees Koedijk**
Clemens Kool***
Franz Palm***

REFERENCES

- Bain, L.J. and M. Engelhardt (1987), 'Introduction to Probability Theory and Mathematical Statistics,' Duxbury, Press Boston.
- Beirlant, J., P. Vynckier and J. Teugels (1996), 'Tail-index Estimation, Pareto Quantile Plots and Regression Diagnostics,' *Journal of the American Statistical Association*, 91, pp. 1659-1667.
- Boothe, P. and D. Glassmann (1987), 'The Statistical Distribution of Exchange Rates: Empirical Evidence and Economic Implications,' *Journal of International Economics*, 22, pp. 297-320.
- Hill, B. (1975), 'A Simple General Approach to Inference about the Tail of a Distribution,' *Annals of Mathematical Statistics*, 3, pp. 1163-1174.
- Huisman, R., K.G. Koedijk, C.J.M. Kool, and F. Palm (2001), 'Tail-index Estimates in Small Samples,' *Journal of Business and Economic Statistics*, 19, pp. 208-216.

* Erasmus University, Rotterdam; ** Erasmus University, Rotterdam, and Maastricht University; *** Maastricht University. The authors would like to thank seminar participants at Maastricht University for their helpful comments. All remaining errors pertain to the authors.

- Kearns, P. and A. Pagan (1997), 'Estimating the Density Tail Index for Financial Time Series,' *The Review of Economics and Statistics*, 79, pp. 171-175.
- Koedijk, K.G. and C.J.M. Kool (1994), 'Tail Estimates and the EMS Target Zone,' *Review of International Economics*, 2, pp. 153-165.
- Koedijk, K.G., P.A. Stork, and C.G. de Vries (1992), 'Differences between Foreign Exchange Rate Regimes: The View from Tails,' *Journal of International Money and Finance*, 11, pp. 462-473.
- Loretan, M. and P.C.B. Phillips (1994), 'Testing the Covariance Stationarity of Heavy-Tailed Time Series: An Overview of the Theory with Applications to Several Financial Data Series,' *Journal of Empirical Finance*, 1, pp. 211-248.
- Lucas, A. (1997), 'Discussion on Fat Tails in Small Samples,' presented at the 5th workshop on Finance and Econometrics in Brussels.
- McCulloch, J.H. (1997), 'Measuring Tail Thickness to Estimate the Stable Index α : A Critique,' *Journal of Business Economics and Statistics*, 15, pp. 74-81.
- Pictet, O., M. Dacarogna and U. Müller (1996), 'Hill, Bootstrap and Jackknife Estimators for Heavy Tails,' *working paper*, Olsen and Associates, Zürich.
- Rogalski, R.J. and J.D. Vinso (1978), 'Empirical Properties of Foreign Exchange Rates,' *Journal of International Business Studies*, 9, pp. 69-79.
- Westerfield, J.M. (1977), 'An Examination of Foreign Exchange Risk under Fixed and Floating Rate Regimes,' *Journal of International Economics*, 7, pp. 181-200.