Journal of International Economics 13 (1982) 171–186. North-Holland Publishing Company

SHORT-RUN FLUCTUATIONS IN FOREIGN EXCHANGE RATES

Evidence from the data 1973–79

Daniel FRIEDMAN

UCLA, Los Angeles, CA 90024, USA

Stoddard VANDERSTEEL

Data Resources, Inc., San Francisco, CA 94104, USA

Received January 1981, revised version received July 1981

The paper examines the statistical properties of daily changes in foreign exchange (FX) rates for nine currencies. It finds that these changes are leptokurtotic, i.e. have long-tailed and sharp peaked histograms. The evidence suggests that this leptokurtosis does not arise from an underlying Paretian stable distribution, nor from a stationary mixture of normal distributions, but rather from normal processes with time-varying parameters. This finding casts doubt on the validity of many standard techniques (t-stats; ARIMA methods). Some moving statistics for estimating the time varying parameters are also presented.

1. Introduction

This paper examines six recent years of daily foreign exchange spot rate movements for nine major currencies. The data are of immediate practical significance to international economists, especially those who analyze risk factors. Econometricians and statisticians may also find the data interesting for methodological reasons. Our analysis indicates that many of the standard statistical procedures based upon assumptions of stationarity and normality are inapplicable to these data. On the other hand, the time series are very long, containing over 1600 observations for each currency. As a result, the data present an unusual opportunity for using non-standard techniques to uncover the underlying patterns.

As we examine the data, we have in mind several questions. To what extent do the currency movements exhibit trends? How can one assess the volatility (i.e. departures from trend) of the various currencies? What underlying economic or institutional processes are consistent with the data? What sorts of behavior by various participants in the FX markets are appropriate in light of the observed fluctuations? This paper does not

0022–1996/82/0000–0000/$02.75

provide definitive answers to any of these questions, but we believe that it sheds considerable light on the first two, and will provide useful information for anyone seeking answers to the second two.

Our emphasis in this paper, then, is more on what can be learned from the data than on methodology. Such econometric novelties as we have employed (e.g. the 'fat tail' test of section 3 and the multi-period variance estimates of section 5) are discussed in more detail in our working paper [Friedman–Vandersteel (1980)], abbreviated henceforth as FV. A much more extensive set of graphs and tables than could be included here can also be found in FV.

The next section provides a basic description of the data, summary statistics and histogram graphs. Sections 3 and 4 examine leptokurtosis, a striking departure from normality exhibited by the data, and test several explanations of its origin. Section 5 introduces and illustrates the use of some moving statistics which seem appropriate for our kind of data. The final section summarizes our conclusions and suggests some interpretations, implications and directions for further work.

2. The data

Our raw data consists of daily foreign exchange (FX) spot rates, quoted in terms of the U.S. dollar, for the currencies listed in table 1, covering the period 1 June 1973 to 14 September 1979. We are primarily concerned with the day-to-day fluctuations in the spot rates, defined $r_t = \log(S_{t+1}/S_t)$, the continuously compounded rate of change of the spot rate on the tth trading day, S_t, to that on the $(t+1)$th trading day, S_{t+1}. (The reader should bear in mind that these rates are correlated across currencies, so the nine series are not independent.)

Table 1 lists summary statistics for the r_t's of the various currencies. Although the mean rates of change are all quite small, ranging from about 0.04 percent per trading day appreciation for the Swiss franc to 0.02 percent depreciation for the Italian lira, very substantial daily changes do sometimes occur, e.g. a 6.7 percent appreciation of the French franc on 4 March 1973 and a 6.7 percent depreciation of the Deutschemark on 1 November 1978. As a first indication of the magnitude of typical fluctuations, we list standard deviations. By this criterion most of our currencies appear to have about the same volatility, roughly 0.5 percent per trading day, the main exception being the Canadian dollar which seems more stable at 0.21 percent. (To obtain annualized rates of appreciation, one multiplies by 261, the usual number of trading days in a year.)

The coefficients of skewness and kurtosis listed in the last two columns of table 1 pertain to the shape of the distribution of the FX rate fluctuations. A preliminary, more qualitative look at a typical distribution is provided in fig.

Table 1
Summary statistics for daily FX rate fluctuations, 1 June 1973–14 Sept. 1979.

Currency	NOBS	Max. (%)	Min. (%)	Mean (%)	Std. Dev. (%)	Skew	Kurt.
German mark	1640	2.92	−6.71	0.024	0.57	−1.18	19.2
Swiss franc	1640	4.41	−6.97	0.039	0.77	−1.07	16.7
Pound sterling	1640	3.96	−4.77	0.010	0.50	−1.22	18.7
Japanese yen	1640	2.76	−6.21	0.010	0.48	−1.62	30.8
Dutch guilder	1640	5.01	−6.82	0.021	0.57	−0.95	25.6
French franc	1640	6.73	−5.00	0.004	0.53	−0.13	26.3
Canadian dollar	1640	0.99	−1.80	−0.010	0.21	−0.58	8.6
Belgian franc	1640	4.16	−5.78	0.017	0.58	−1.10	18.4
Italian lira	1640	4.14	−5.54	−0.021	0.49	−0.93	27.6

Data source: DRI-FACS data base, N.Y. daily opening asking rates for spot FX, Data Resources, Inc.

1. This histogram indicates the probability density of the observed fluctuations for the deutschemark. The superimposed curve indicates the normal density of the same mean and standard deviation. All the currencies give rise to essentially unimodal, approximately symmetric (i.e. 'bell-shaped') distributions, but with some anomalies, to which we now turn.

The coefficient of skewness (defined as $SK=E(r_t-\mu)^3/\sigma^3$, the third moment around the mean divided by the cube of the standard deviation) detects asymmetries; it is zero for a symmetric distribution. A glance at table 1 indicates moderately negative SK for all currencies with the exception of the essentially unskewed French franc. Such negative skewness suggests that the

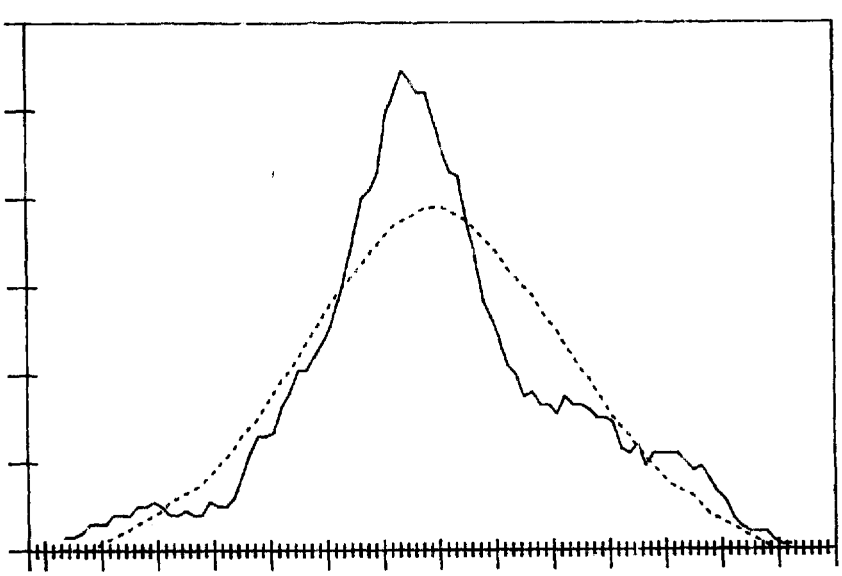

Fig. 1. Distribution of German mark exchange rates (line: observed; dot: theoretical normal).

lower tails of the distribution are longer than the upper, i.e. the largest downward fluctuations outweigh the largest upward fluctuations, a suggestion confirmed by comparing columns 2 and 3 of the table. Skewness will be discussed again in section 5; for now we will comment only that much of the observed skewness can be attributed to the events of a single trading day (for most currencies, 1 November 1978, a day marked by massive intervention by the Fed in support of the U.S. dollar), so it is not unreasonable for most purposes to regard the distributions as symmetric.

3. Leptokurtosis

The last column of table 1 suggests a more serious anomaly. The coefficient of kurtosis (defined as $KURT = E(r_t - \mu)^4/\sigma^4$, the fourth moment around the mean divided by the fourth power of the standard deviation) indicates how the probability 'mass' is spread in the distribution. Kurtosis is lower the more of the mass of the probability distribution is concentrated near $\mu \pm \sigma$. A uniform distribution, for instance, has low kurtosis as does any distribution with a modest mode and short tails. (A 'low' kurtosis coefficient is one between 1 (the minimum value) and 3, the value for a normal distribution.) The observed distributions of FX rate fluctuations all have kurtosis coefficients considerably larger than 3, a condition known as *leptokurtosis*.[1] The basic symptoms of leptokurtosis, more massive tails and a sharper central peak than normal, appear in all the histograms.

Many standard statistical procedures (e.g. the usual significance tests involving t-statistics) become invalid in the presence of significant leptokurtosis, so it is important to confirm the diagnosis. The coefficient of kurtosis for a sample is not a reliable indicator because a few extreme outliers can very quickly raise the kurtosis coefficient in even fairly large samples. One rather direct test for non-normal tails is to examine the number of observations lying in the tails, say at more than three standard deviations from the mean. In large samples drawn from a normal distribution, approximately 0.26 percent (or about 4 out of 1600) of the observations should lie at this distance. In samples of 1640 observations, we find from 20 (for the Dutch guilder) to 36 (for the Swiss franc) in the 3σ tails; the probabilities of drawing such samples from normal distributions are only 0.0000026 percent and $< 10^{-20}$, respectively, confirming that the tails are indeed more massive than normal. See FV for a discussion of this test.

For detecting abnormally long tails, Fama and Roll (1971) recommend highly the studentized range statistic, SR = (largest observation − smallest observation)/s, where s = sample standard deviation. Casual interpolation of the SR tables in David, Hartley and Pearson (1954) reveals for a sample of our size that any value of SR in excess of 8.5 indicates longer-than-normal

[1] Not to be confused with *leptospirosis*, a disease of humans and domesticated animals caused by spirochetes of the genus leptospira, or *leptonecrosis*, decay of the phloem tissues in plants.

tails at the 0.005 confidence level. From table 1 one can readily compute that *SR* ranges from 13 for the Canadian dollar to 21 for the French franc, once again abundantly confirming leptokurtosis.

4. The origin of leptokurtosis

The literature [see Granger and Orr (1972), for instance] provides two competing explanations for leptokurtosis. In the first, call it H_0, leptokurtosis is an irreducible property of the data; the time series is assumed to arise by drawing observations from a distribution which is *Paretian symmetric stable* (PSS). These distributions form a three-parameter family, and are specified by their log characteristic functions $\log \phi(t) = idt - |ct|^{\alpha}$, where d is a location (i.e. trend) parameter, c is a scale (volatility) parameter, and α is the exponent.[2] These distributions are called 'stable' because they are well-behaved under summing and the limiting operations of the central limit theorem. For instance, a sum of independent PSS variables of exponent α is also PSS of exponent α.

The stability property is desirable in modeling economic time series such as FX rate fluctuations which can be regarded as sums of independent random events, and the only symmetric leptokurtotic distributions with the stability property belong to the PSS family and have $\alpha < 2$. Unfortunately all such distributions have infinite variance. Of course, the variance of any finite sample drawn from one of these distributions will be finite, but will become unbounded as the sample size increases.

Another explanation of leptokurtosis is that each sample is drawn from two or more different normal distributions, i.e. a sample arises from a mixture of normal distributions with different means and variances. Such a hypothesis avoids the problems of infinite variance (i.e. divergent sample variance) but abandons the stability property. We can distinguish two versions of this explanation. In the first — call it H_1 — the mixture of normal distributions from which an observation is drawn is always the same. In the second — call it H_2 — we draw observations from a normal distribution whose parameters are time-dependent.[3] The explanations, H_0–H_2, of leptokurtosis are not exhaustive, but appear to be the most attractive and the most commonly mentioned in the literature.

Several criteria have been suggested for choosing among the explanations of leptokurtosis. Perhaps the first to come to mind is to take increasing

[2]If $\alpha = 2$, the distribution is normal with mean d and standard deviation $c/\sqrt{2}$. If $\alpha = 1$, the distribution is Cauchy centered at d with semi-interquartile range c. If $1 < \alpha < 2$, the stable distribution is leptokurtotic and has no known elementary expression for its density or cumulative distribution function.

[3]To make these explanations more explicit, assume $r_t \sim N(\mu_t, \sigma_t)$. Under H_1, the parameters (μ_t, σ_t) are independent and identically distributed; in particular the distribution is time-independent, so we can refer to this as the autonomous case. Under H_2, (μ_t, σ_t) are non-trivial functions of time and consequently the r_t series is non-stationary.

samples $R_k=\{r_1,r_2,..,r_k\}$, $k\leqq N(=\text{NOBS})$ and to plot the sample variances s_k^2 against k; if the s_k^2 appear to diverge, H_0 would be supported. The trouble with this procedure is that one cannot really say if a series diverges by examining its initial segment. Also, some forms of H_0 are compatible with apparent divergence of s_k^2.

A more promising test is based on the fact that if independent observations are drawn from a stable distribution with scale parameter c_1 and exponent α, then the stability property implies that their sum has scale parameter $c_k=c_1k^{1/\alpha}$ (and exponent α). Therefore, α can be estimated by regressing $\ln c_k$ on $\ln k$ for various values of k; the slope coefficient should be $1/\alpha$ and the intercept $\ln c_1$. See Fama and Roll (1971, p. 334) for background and FV for details of this procedure. Table 2 reports the estimates[4] of c_k and α, which strongly suggest $\alpha=2$, thus supporting H_1 or H_2 at the expense of H_0. (Recall H_0 asserts $\alpha<2$.)

Table 2
Estimates of k-period standard deviation S_K and the characteristic exponent α.

Currency	S_1 $\times 10^{-2}$	S_2	S_5	S_{11}	S_{21}	S_{65}	S_{131}	S_{261}	α	Std. error
WG	0.573	0.838	1.304	1.881	2.633	4.590	6.650	8.743	2.032	0.004
SW	0.772	1.129	1.743	2.516	3.540	6.253	9.185	12.243	2.004	0.004
UK	0.495	0.711	1.096	1.642	2.237	4.067	5.654	7.806	2.012	0.003
JA	0.478	0.686	1.092	1.577	2.155	3.722	5.571	7.176	2.040	0.005
DU	0.567	0.820	1.274	1.863	2.618	4.546	6.450	8.590	2.037	0.003
FR	0.527	0.777	1.205	1.761	2.446	4.276	6.270	8.186	2.011	0.005
CA	0.214	0.306	0.472	0.713	0.987	1.749	2.418	3.618	1.984	0.003
BE	0.580	0.844	1.302	1.865	2.712	4.555	6.717	8.378	2.053	0.007
IT	0.493	0.710	1.114	1.616	2.280	3.912	5.796	7.996	2.000	0.003
N0	0.579	0.825	1.303	1.901	2.618	4.757	6.575	9.502	1.995	0.002
N1	0.594	0.848	1.3[illegible]4	1.925	2.709	4.895	6.813	9.632	1.995	0.003
N2	0.592	0.834	1.3[illegible]3	1.958	2.686	4.692	6.805	9.546	2.003	0.002

Notes: N0–N2 are 'artificial currencies' described in section 5. S_k is an estimate of the scale parameter c_k. Std. error is the standard error of the slope coefficient in the regression of $\log S_k$ on $\log k$.

The test most favored by Fama and Roll is similar in spirit, but based on order statistics. They first show that for $1\leqq\alpha\leqq 2$ the scale parameter $c(=c_1)$ can be efficiently estimated by $\hat{c}=\frac{1}{2}(\hat{X}_{0.72}-\hat{X}_{0.28})/0.827$, where $\hat{X}_f$ is the $(N+1)f$ order statistic (i.e. the value at the $100f$ percentile in the observed sample). Then the statistic

$$\hat{Z}_{0.97}=(\hat{X}_{0.97}-\hat{X}_{0.03})/2\hat{c}=0.827(\hat{X}_{0.97}-\hat{X}_{0.03})/(\hat{X}_{0.72}-\hat{X}_{0.28})$$

[4]Table 2 uses the k-period standard deviation s_k to estimate scale. Although the scale estimator $\hat{c}_k$ (described in the following paragraph of the text) might be superior for present purposes, the s_k estimates are more germane for issues discussed in section 5.

is an estimator of the 0.97 fractile (97th percentile) of the standardized ($d=0$, $c=1$) PSS distribution of exponent α. There is a monotonic (decreasing) relationship between $\hat{Z}_{0.97}$ and α, tabulated in Fama–Roll (1968, p. 822). Thus, one can form the statistic $\hat{Z}_{0.97}$ and use the table to derive $\hat{\alpha}$ from any sample. If one forms a new sample of size $\simeq N/k$ by aggregating groups of k successive observations in the original sample, the resulting estimates $\hat{\alpha}_k$ should be essentially independent[5] of the degree of aggregation k if H_0 is correct. On the other hand, if H_1 is correct, $\hat{\alpha}_k$ should approach 2 as k increases, since the (aggregated) observations become more nearly 'identically distributed' and the variance is finite. H_2 implies slightly different behavior: even for large k, the (aggregate) observations may not have nearly identical distributions, so $\hat{\alpha}_k$ may remain well below 2. However, if one aggregates *randomly selected* but non-overlapping groups of k individual observations, rather than *successive* observations, then H_2 becomes indistinguishable from H_1 and therefore implies convergence of α_k to 2.

Table 3 provides strong support for H_2 over both H_0 and H_1. The column $k=1$ lists full-sample estimates of the exponent α, given that H_0 is correct; the extent to which an entry is less than 2 may be taken as an index of leptokurtosis for the observed distributions. The stability property of PSS distributions implies that the entries of the other columns also are estimates of the same α if H_0 is correct, although the error will increase with k since the sample size $[N/k]$ decreases. However, for each currency the entry for $k=20$ exceeds the $k=1$ entry, casting severe doubt on H_0. For comparative purposes we generated 1640 independent normally distributed random numbers with the mean and variance of the deutschemark series and applied the α-estimating procedure to them, with results listed in row 10. All these entries are estimates of $\alpha=2$; there appears to be no tendency for them to increase or decrease with k.

It is not clear that the estimates of α are converging to 2 in table 3, part A, but such convergence becomes quite plausible in part B (if one bears in mind the error increases with k). In the latter, the dates t were randomly permuted[6] before aggregation. Under H_1 this should have no systematic effect on the estimates of α; in actuality the estimates of α for $k=20$ are higher in part B of table 3 than in part A, for every currency except the French franc (1.74 vs. 1.72); even in this case, one gets a stronger impression of convergence to 2 in part B than part A.

[5]Because of the stability property. Actually, as Fama and Roll (1971) point out, a slight bias in the tables (due to rounding off) would cause the $\hat{\alpha}$ to drift downward slightly as k increases.

[6]Actually, five different random permutations were used and the results averaged in order to reduce errors. The informal arguments of the text are not affected, but if one wished to perform quantitative tests on the table entries, this averaging should be taken into account. The separate results for each random permutation (as well as results for replacing 97 by 95 and 99 in the α-estimator) are included in FV; all tell basically the same story.

Table 3
Estimates of the exponent α.

$k=$	1	2	5	10	20
A. *Direct* aggregation of degree k					
1. German mark	1.45	1.47	1.50	1.46	1.63
2. Swiss franc	1.39	1.46	1.36	1.60	1.63
3. British pound	1.29	1.39	1.39	1.54	1.62
4. Japanese yen	1.11	1.18	1.30	1.37	1.33
5. Dutch guilder	1.48	1.47	1.47	1.57	1.49
6. French franc	1.36	1.38	1.37	1.54	1.74
7. Canadian dollar	1.55	1.58	1.67	1.48	1.74
8. Belgian franc	1.39	1.39	1.45	1.54	1.52
9. Italian lira	1.12	1.23	1.14	1.18	1.24
10. Random DM	2.12	1.96	2.12	2.44	2.00
B. *Scrambled* aggregation of degree k (average of five random permutations)					
1. German mark	1.45	1.58	1.72	1.74	1.92
2. Swiss franc	1.39	1.48	1.65	1.77	1.72
3. British pound	1.29	1.47	1.64	1.75	1.76
4. Japanese yen	1.11	1.39	1.65	1.69	1.75
5. Dutch guilder	1.48	1.58	1.79	1.76	1.79
6. French franc	1.36	1.48	1.71	1.86	1.72
7. Canadian dollar	1.55	1.70	1.84	1.88	2.06
8. Belgian franc	1.39	1.50	1.63	1.93	2.15
9. Italian lira	1.12	1.37	1.47	1.54	1.67
10. Random DM	2.12	1.94	1.97	1.98	1.78
11. 25 Monte Carlo trials	2.00 ± 0.08	1.99 ± 0.06	2.02 ± 0.07	2.02 ± 0.10	2.00 ± 0.17

$f=0.97$; NOBS$=[1640/k]$.

In order to check this impression, we performed the following Monte-Carlo experiment, suggested by an anonymous referee. We generated twenty-five random deutschemark series, each time applying the α-estimating procedure to the series after randomly permuting the dates. The mean outcome of the twenty-five trials is reported in line 11 of table 3, part B, with the standard deviation of the trial outcomes indicated underneath. The population means should all be 2.00 by construction, but the standard deviations should decrease with k, as noted above. One should interpret the entry of 0.17 in the $k=20$ column, for instance, as an indication of the standard error on the $\hat{\alpha}_{20}$. By this measure, $\hat{\alpha}_{20}$ has 'converged' to within two standard errors of 2.00 for all nine currencies in table 3, part B, but for only two currencies (the French franc and the Canadian dollar) in part A.

The virtual unanimity among currencies in demonstrating that convergence to 2 is improved by scrambling dates before aggregating leads us to conclude that H_2 provides a better explanation of the data than H_1 (and, a fortiori, better than H_0).

5. Treatment of leptokurtotic data

Granted that our data suffers from severe leptokurtosis, and that this condition may arise from an underlying normal process whose parameters vary over time, how does one proceed? In particular, how can one recover the time-paths of the parameters?

It is necessary at this point to make more specific assumptions about the nature of the time dependence. If r_t is drawn from an arbitrary normal distribution $N(\mu_t, \sigma_t)$, then virtually anything is possible unless some restrictions are placed on the μ_t, σ_t. One reasonable hypothesis is that the parameters vary *slowly*; say changes are usually slight within a quarter. One can then use moving statistics to estimate μ_t and σ_t.

The most direct approach would be to select some 'window length' — say one quarter, about 65 trading days — and estimate μ_t and σ_t by the mean and standard deviation of $R_t^{65}=\{r_{t-64}, r_{t-63}, \ldots, r_t\}$, the 65 most recent observations. However, these estimators are neither 'robust' nor 'resistant'.[7] For instance, it will take most of the quarter before they reflect a shift in the parameters, and an occasional outlier will induce a bias in the estimators that will persist for the same length of time. The second type of problem would be mitigated by taking a longer 'window' than 65, but then shifts of μ_t and σ_t become increasingly hard to detect. One could avoid these problems, at the cost of a major loss of efficiency, by using the median instead of the mean of R_t^{65} to estimate $\hat{\mu}_t$, and an appropriate multiple of the difference between fractiles of R_t^{65} to estimate $\hat{\sigma}_t$.

Our prescription, adapted from Cleveland and Kleiner (1975), is a compromise between the last two approaches. Keeping the window length at 65, one defines the moving mid-mean MMM_t as the mean of the central 50 percent of the observations R_t^{65}. That is, one throws out the largest 25 percent and the smallest 25 percent of the observations (thus gaining 'resistance' to outliers and speeding adjustment to shifts of μ_t) and then takes the mean of what remains (thus maintaining reasonable efficiency). One similarly defines the moving upper (lower) semi-mid-means UMM_t (LMM_t), essentially by taking the MMM of the upper (lower) 50 percent of R_t^{65} — see the FV for details. UMM_t (LMM_t) is a robust, resistant and reasonably efficient estimator of essentially the 75th (25th) percentile point.

Fig. 2 plots UMM_t, MMM_t, and LMM_t, $t=66, 70, 75, \ldots, 1640$, for the deutschemark. The middle line (MMM_t) indicates the trend, while the outer lines (UMM_t, LMM_t) enclose a 50 percent confidence interval around the trend. One observes a trend which reverses itself several times during 1973–75 before stabilizing near 0 during 1976–77. Meanwhile, the DM became

[7]Roughly speaking, an estimator is 'robust' if its performance is relatively insensitive to small changes in the distribution, and 'resistant' if it is relatively insensitive to changes of a small part of the data. See Mosteller and Tukey (1977) for a discussion of these concepts, and Krasker (1980) for recent applications to regression models.

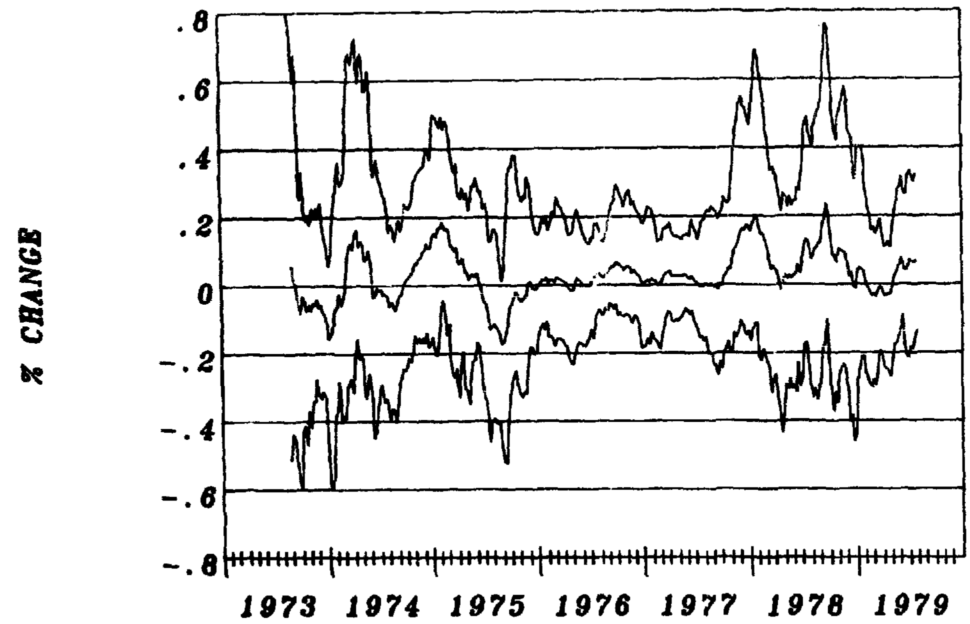

Fig. 2. German mark, mid-means.

increasingly less volatile, indicated by the narrowing gap between UMM_t and LMM_t from 1973 to 1977. In late 1977, the DM began a new upward trend, associated with much higher volatility, which (apart from a short lull in early 1978) persisted until almost the end of 1978; from then until the end of our sample period (September 1979) the DM was relatively trendless and less volatile.

The patterns displayed in the mid-mean graphs for the other currencies (see FV) are quite varied, but equally striking. Are they perhaps just artifacts of our statistical techniques and of no economic significance? The significance of the patterns can be confirmed qualitatively in two different ways. First, one can in effect use an 'experimental control' by applying the mid-mean treatment to random noise.[8] We randomly drew 1640 observations from a normal distribution of mean 0.00024 and standard deviation 0.0058, the values for the deutschemark. Fig. 3 shows the resulting MM graph. One detects apparent short-lived trends, but volatility does not appear to change, and the qualitative difference between the MM graph for the 'random German mark' and those for real currencies is quite striking to most observers. The second argument for the significance of the MM graphs, although even more 'qualitative', may be more persuasive to the practical-minded reader; in almost every case, the major observed patterns can be

[8]We wish to thank Ed Leamer for this suggestion. We actually performed the experiment described in the next sentence of the text several times, all resulting in graphs closely resembling fig. 3.

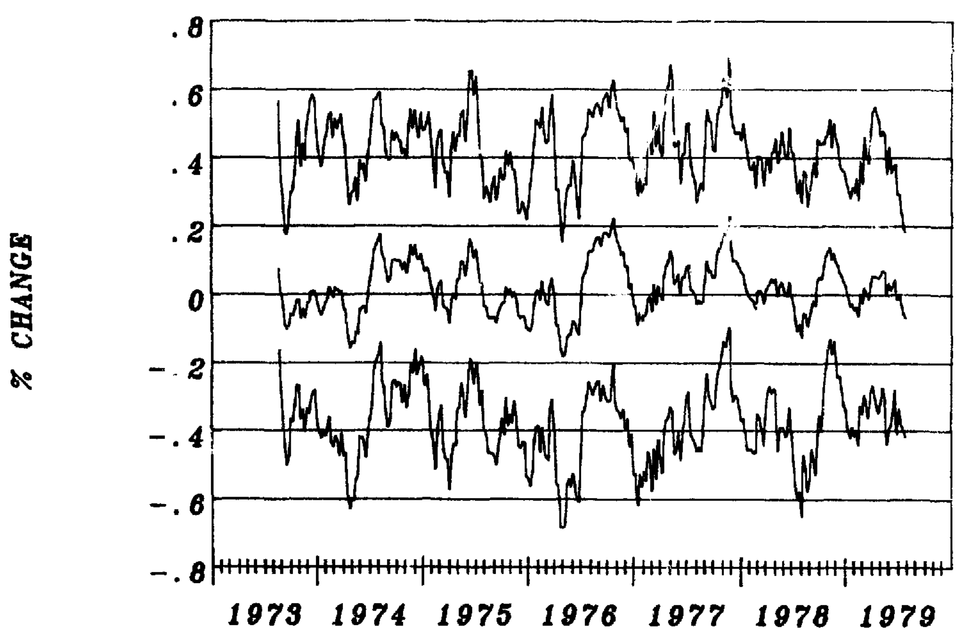

Fig. 3. Random German mark, mid-means.

closely associated with historically verifiable developments in the FX markets.[9]

Supposing that MMM_t is a reasonable estimator of μ_t, we still need to estimate σ_t. Since UMM_t (LMM_t) estimates the 0.75 (0.25) fractile, which should be near $\mu_t+0.693\sigma_t$ ($\mu_t-0.693\sigma_t$), it seems reasonable to set $\hat{\sigma}_t=\frac{1}{2}(UMM_t-LMM_t)/0.693$. A full explanation of $\hat{\sigma}_t$ is provided in FV. Fig. 4 graphs the resulting volatility estimates for the deutschemark.

The trio of moving statistics also provides a way of detecting asymmetries, more robust than the skewness coefficient. One compares the difference between the 50th and 75th percentiles to the difference between the 25th and 50th. If the former is larger (smaller), the distribution is positively (negatively) skewed. Thus, the moving statistics $SKEW_t=(UMM_t-MMM_t)-(MMM_t-LMM_t)=UMM_t+LMM_t-2MMM_t$ indicates asymmetries. In almost all cases, $SKEW$ looks like white noise, justifying our early assumption of symmetry. Fig. 5 for the DM is typical.

[9]For instance, the 1973–75 period marked the emergence of the 'managed float' regime in FX markets, and was characterized by widening inflation rate differentials between the U.S. and West Germany. The latter part of 1977 and first part of 1978 saw the inflation differentials widen again and the U.S. current account move sharply into deficit.

When we have shown and explained our MM graphs to FX dealers and economists, we have heard a wide variety of interpretations of the observed patterns — indeed, they could perhaps even serve as a Rorschach test for eliciting a subject's basic approach to the determination of FX rates. We have been encouraged, however, by the ability of many expert subjects to correctly identify the currency from the pattern, and by the ability of all subjects to distinguish artificial currencies (the 'random DMs') from real currencies.

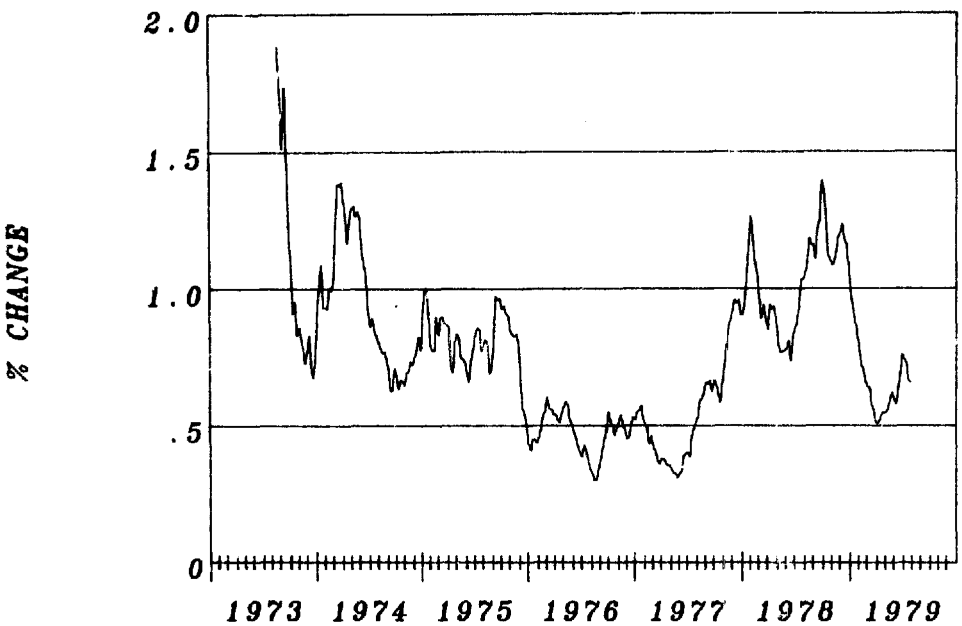

Fig. 4. Volatility of the German mark exchange rate.

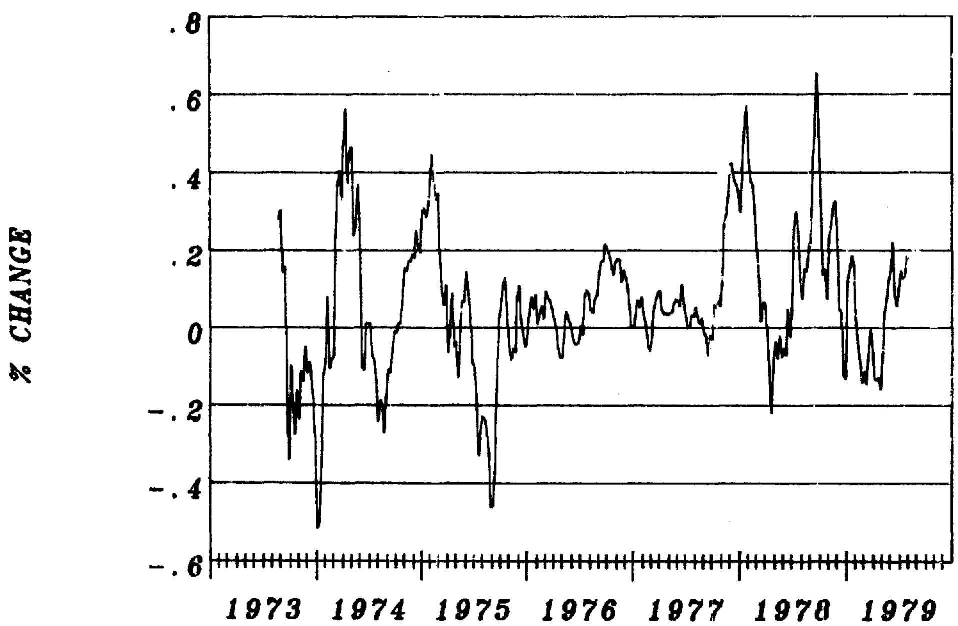

Fig. 5. Skewness of the German mark (upper + lower − 2 mid-mean).

Another possible complication, heretofore ignored but still requiring examination, is autocorrelation. A discussion in Granger and Orr (1972) suggests that the usual autocorrelation coefficients provide useful information even in the presence of leptokurtosis. Accordingly, we computed the autocorrelation coefficients $\rho_s = \text{corr}(r_t, r_{t-s})$ for various lags s and all currencies; the results are summarized in table 4. Most coefficients are

Table 4

Partial autocorrelation coefficients for daily FX fluctuations, 1 June 1973–14 Sept. 1979.

	Lag								
Currency	1	2	3	4	5	10	15	16	20
German mark	0.022 (0.91)	−0.082 (−3.31)	0.007 (0.30)	0.016 (0.63)	−0.059 (−2.40)	0.026 (1.05)	0.050 (2.01)	−0.004 (−0.18)	0.028 (−1.14)
Swiss franc	−0.001 (−0.06)	−0.106 (−4.28)	0.004 (0.18)	−0.017 (−0.71)	−0.041 (−1.66)	0.015 (0.61)	0.079 (3.20)	0.004 (0.14)	−0.062 (−2.52)
Belgian franc	−0.026 (−1.06)	−0.095 (−3.83)	0.016 (0.65)	0.041 (1.67)	−0.024 (−0.97)	0.034 (1.38)	0.011 (0.46)	−0.001 (−0.05)	0.015 (0.60)
Italian lira	0.123 (4.98)	−0.007 (−0.27)	−0.061 (−2.48)	−0.55 (−2.24)	0.001 (0.04)	−0.028 (−1.15)	0.070 (2.82)	0.092 (3.72)	0.042 (−1.71)
Dutch guilder	−0.005 (−0.21)	−0.064 (−2.58)	0.004 (0.16)	0.023 (0.94)	−0.021 (−0.86)	0.059 (2.38)	0.021 (0.86)	0.004 (0.15)	−0.021 (−0.86)
French franc	0.034 (1.38)	−0.044 (−1.78)	0.006 (0.23)	0.026 (1.06)	−0.024 (−0.98)	0.028 (1.15)	0.026 (1.05)	−0.006 (−0.24)	−0.013 (−0.53)
Pound sterling	0.008 (0.34)	−0.038 (−1.56)	0.022 (0.90)	−0.026 (−1.06)	0.022 (0.91)	0.092 (3.71)	0.046 (1.85)	0.013 (0.51)	−0.028 (−1.15)
Japanese yen	0.067 (2.71)	−0.015 (−0.62)	0.046 (1.88)	0.015 (0.61)	−0.038 (−1.52)	0.067 (2.71)	0.052 (2.10)	−0.026 (−1.04)	0.003 (0.14)
Canadian dollar	0.004 (0.15)	−0.024 (−0.95)	0.039 (1.57)	0.030 (1.21)	0.040 (1.60)	0.035 (1.42)	−0.031 (−1.24)	−0.034 (−1.38)	0.001 (0.04)

Note: t-statistics in parentheses.

insignificant, but there are more 'significant' t-statistics than one would expect from uncorrelated normal data. Bearing in mind that the data are leptokurtotic, not normal, these t-statistics must not be taken too seriously. In order to check the robustness and stability of some of these 'significant' coefficients, we examine more closely in table 5 the four largest: the 1st and 16th order autocorrelations of the Italian lira, 10th UK pound and the 2nd Swiss franc. The first of these appears to be the largest and most significant, but evidently the observed autocorrelation arises entirely from events in the first half of the sample period, since the coefficient changes sign and becomes insignificant in the second half of the sample. The second line shows that this coefficient again becomes insignificant when the sample is trimmed by throwing out the largest and the smallest 10 percent of the observations.[10] Evidently this full sample estimate of 0.123 arises mostly from a few extreme observations, probably during the precipitous depreciation of the lira in early 1976.

A similar pattern appears for the other large autocorrelation coefficients: estimates become much less significantly different from zero in the trimmed sample, and differ between the two sample periods. We conclude that the FX rate fluctuations are probably not autocorrelated to any significant degree, that the larger observed autocorrelation coefficients are generally not stable, and probably arise mostly from 'random' placement of the extreme fluctuations.

6. Summary and concluding remarks

Our analysis of daily fluctuations in FX rates shows quite clearly that they are not normally distributed. Very large fluctuations (and also very small fluctuations) are more common than one might expect. In the presence of such leptokurtosis, one must be careful in applying analytic techniques which presume normality.

As a first step in dealing with the malady of leptokurtosis, we examined several possible explanations of its origin. We rejected the hypothesis H_0 that the leptokurtosis is irreducible and due to an underlying process of infinite variance. We also rejected the hypothesis H_1 that leptokurtosis is due to an autonomous mixture of normal processes. (A priori, H_1 has a reasonable institutional interpretation: one might hypothesize that a FX rate fluctuation

[10]In more detail, our procedure was to flag the largest 5 percent and the smallest 5 percent of the r_t's (the 'outliers') and to omit comparisons involving these outliers in computing the autocorrelation coefficients. Specifically, let s_t be standardized r_t's, let F be the set of dates of the outliers, T be the full set of dates, and l be the autocorrelation lag. Then define $E=\{t \in T \mid t \notin F$ and $t-l \notin F\}$; in line 2 of table 5 we report $(\#E-1)^{-1}\sum_{t\in E} s_t s_{t-l}$. It turns out that $\#E$ is 80–84 percent of $\#T=1640$ in our four computations.

Table 5
Stability of selected autocorrelation coefficients.

Currency:	Italian lira		Swiss franc	UK pound
Lag:	1	16	2	10
1. Full sample	0.123	0.092	−0.106	0.92
	(4.98)	(3.72)	(−4.28)	(3.71)
2. 10% trimmed sample	0.038	0.066	−0.052	0.030
	(1.41)	(2.46)	(−1.92)	(1.10)
3. First half of sample (2 June 1973–23 Sept. 1976)	0.176	0.094	−0.137	0.076
	(5.04)	(2.70)	(−3.92)	(2.165)
4. Second half of sample (24 July 1976–14 Sept. 1976)	−0.031	0.052	−0.073	0.094
	(−0.88)	(1.482)	(−2.09)	(2.70)

Notes: t-statistics in parentheses. 10% trimmed sample involves estimates based on the central 80% of the ordered data.

for a day marked by significant Central Bank intervention has a different probability distribution than that for a day in which the markets were left to their own devices. If there were no discernible time trend to the interventions, H_1 would be appropriate.)[11] The data appear to support the hypothesis H_2 that there is an underlying normal process which generates the fluctuations, but that the process changes over time. An interpretation is that both the trend *and volatility* of FX rate movements are affected by changing economic and institutional factors.

Under the assumption the process changes slowly. we derived estimates for the trend and volatility from moving statistics with desirable statistical properties. The patterns uncovered by this procedure seem to agree well with known historical developments in the FX markets. Of course, much more careful and quantitative work is needed if these moving statistics are to be relied upon.

Some implications of our analysis for international economists are clear. The correctness of H_2 casts real doubt on the validity of any data analysis employing ARIMA or other techniques which presume the stationarity of the r_t series.[12] If one is interested only in testing the hypotheses about which variables affect exchange rates, the procedures in White (1980) might suffice. Beyond this, our results suggest that an autonomous two-parameter model cannot adequately characterize FX risk for many purposes.

[11]Another plausible argument for H_1, suggested by an anonymous referee, is that 'shocks' reflected in the r_t's are drawn from two distributions, one for each country whose exchange rate is under consideration.

[12]The usual 'fix-up' for non-stationarity — further differencing of the time series — does not appear to help. Empirically, we find that *second* differences of log spot rates for the deutschemark are still leptokurtotic, and a rough α test still seems to favor H_2. Of course, if one believes our model of $r_t \sim N(\mu_t, \sigma_t)$, with μ_t and σ_t varying slowly with time, further differencing has no theoretical justification anyhow.

A time-dependent model, with μ_t and σ_t some simple functions of time, can be made consistent with our analysis of the data and may be appropriate for some theoretical work. Perhaps a more fruitful approach, also consistent with our findings, would be to estimate μ_t and σ_t as functions of time-varying economic and institutional variables. For instance, one could use our estimates $\hat{\mu}_t$ and $\hat{\sigma}_t$ as dependent variables in regressions on economic time series, perhaps augmented with institutional proxy variables. The resulting fitted variables could be of value in empirical studies, particularly those for which risk factors are important.

Certainly more data analysis is called for. For instance, the covariance structure of short-run fluctuations is of great interest, and robust methods would seem desirable. Daily forward FX rates are also available and should be analyzed. Such analysis might well yield important insight into the behavior of investors, importers and exporters, banks and other international transactors.

Acknowledgements

The authors would like to thank Doug Martin, Dennis Sheehan and Ed Leamer for helpful suggestions and guidance to the literature, Guiseppe Dose and Ray Peters of Bank of America for valuable comments on FX market behavior, and two anonymous referees of this journal for suggesting ways to strengthen our paper. Able research assistance was provided by Rudi Schadt and some helpful programming hints by Dan Grafer.

References

Cleveland, William S. and Beat Kleiner, 1975, A graphical technique for enhancing scatterplots with moving statistics, Technometrics 17, 147–454.

David, H.A., H.O. Hartley and E.S. Pearson, 1954, The distribution of the ratio, in a single normal sample, of range to standard deviation, Biometrika 41, 482–493.

Fama, Eugene F. and Richard Roll, 1968, Some properties of symmetric stable distributions, Journal of the American Statistical Association 63, 817–836.

Fama, Eugene F. and Richard Roll, 1971, Parameter estimates for symmetric stable distributions, Journal of the American Statistical Association 66, 331–338.

Friedman, Daniel and Stoddard Vandersteel, 1980, Short-run fluctuations in foreign exchange rates: An exploration of the data, Discussion Paper No. 171, Department of Economics, UCLA, May.

Granger, Clive W.J. and Daniel Orr, 1972, 'Infinite variance' and research strategy in time series analysis, Journal of the American Statistical Association 67, 275–285.

Krasker, William S., 1980, Estimation in linear regression models with disparate data points, Econometrica 48, 1333–1346.

Mosteller, Frederick and John W. Tukey, 1977, Data analysis and regression (Addison-Wesley Publishing Co., Reading, Mass.).

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, Econometrica 48, 817–838.