Chapter 3

Preliminary Data Analysis

In this chapter we present a preliminary analysis of the time series of exchange rates for the pound, yen and deutsche mark relative to the US dollar. This analysis is based mainly on graphical methods, and shows how the standard techniques fail to adequately describe the changing volatility of these data.

The need for a transformation

The data of financial time series is the trivariate daily exchange rate series of the pound sterling, Japanese yen and the German deutsche mark measured against the US dollar, running from 3 January 1986 to 12 April 1994, yielding 2158 observations. Table 3.1 shows the numerical summaries of each daily exchange rate.

Table 3.1: Numerical summaries of exchange rates relating to US dollar

col	variable	size	mean	st dev	min	max
1	pound/dollar	2158	0.606	0.055	0.498	0.728
2	100 yen/dollar	2158	1.359	0.179	1.006	2.042
3	DM/dollar	2158	1.759	0.198	1.389	2.474

Exchange rates: 3 Jan 86 - 12 Apr 94

In Table 3.1, for the Japanese yen, we multiply by 100 against the US dollar, to make this exchange rate more comparable with the others. The result shows that the mean exchange rate and standard deviation of the pound sterling against the US dollar is small compared to the yen and the deutsche mark against the US dollar. The result also shows the mean and standard deviation of exchange rates of the yen and the deutsche mark are closer to each other.

We also plot these data as time series, as shown in Figure 3.1.

Figure 3.1: Daily exchange rates from 3 January 1986 to 12 April 1994

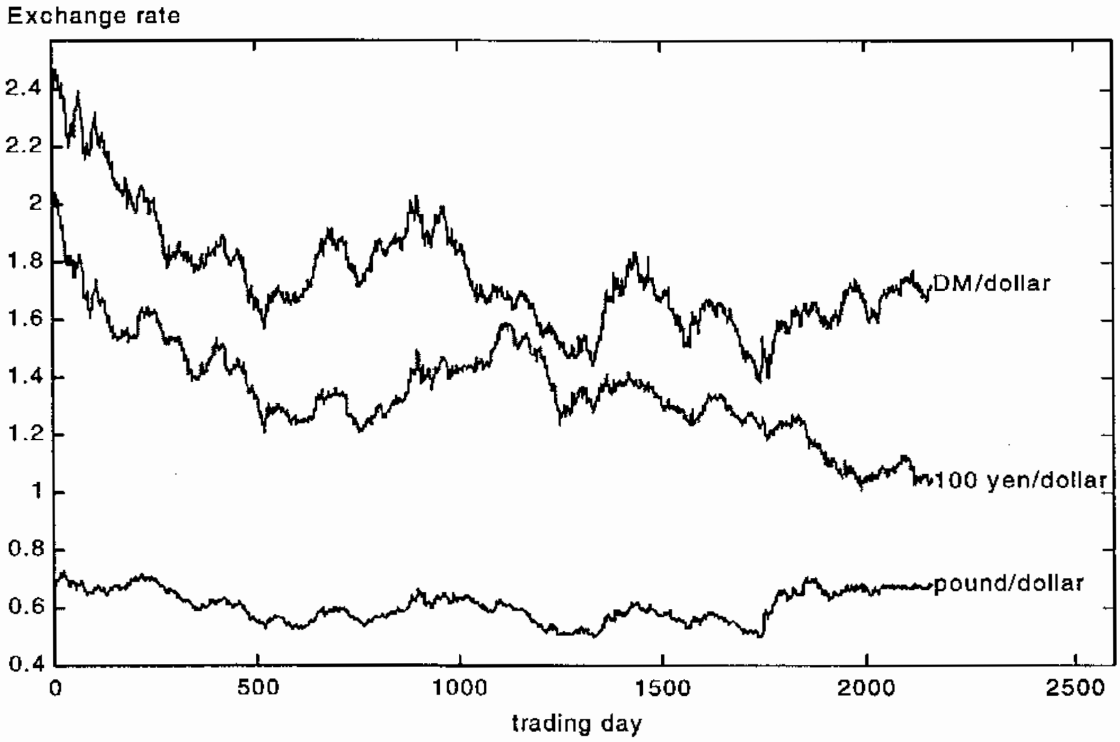

In Figure 3.1, the pound sterling shows a stable series compared to the yen and the deutsche mark. The yen and deutsche mark fluctuate with similar variability. There is a clear change approximately after day 1750 day where the yen decreases but the deutsche mark increases.

Next we form the samples comprising 100 successive trading days, and then look at the relation between the mean and the standard deviation of each sample, shown in Figure 3.2. These data are plotted on a logarithm scale to show the relation more clearly. As the mean increases so does the standard deviation. This suggests a need to transform the data.

Figure 3.3 shows the relation between the standard deviation and the mean of the samples of 100, after taking natural logarithms of the exchange rates. In this plot, the standard deviations (but not the means, which can be negative) are again expressed on a logarithmic scale.

Figure 3.3 shows no correlation between the mean and standard deviation of the transformed data. We conclude that the log transformation is appropriate.

Figure 3.2: Relation between standard deviation and mean in 100 trading days

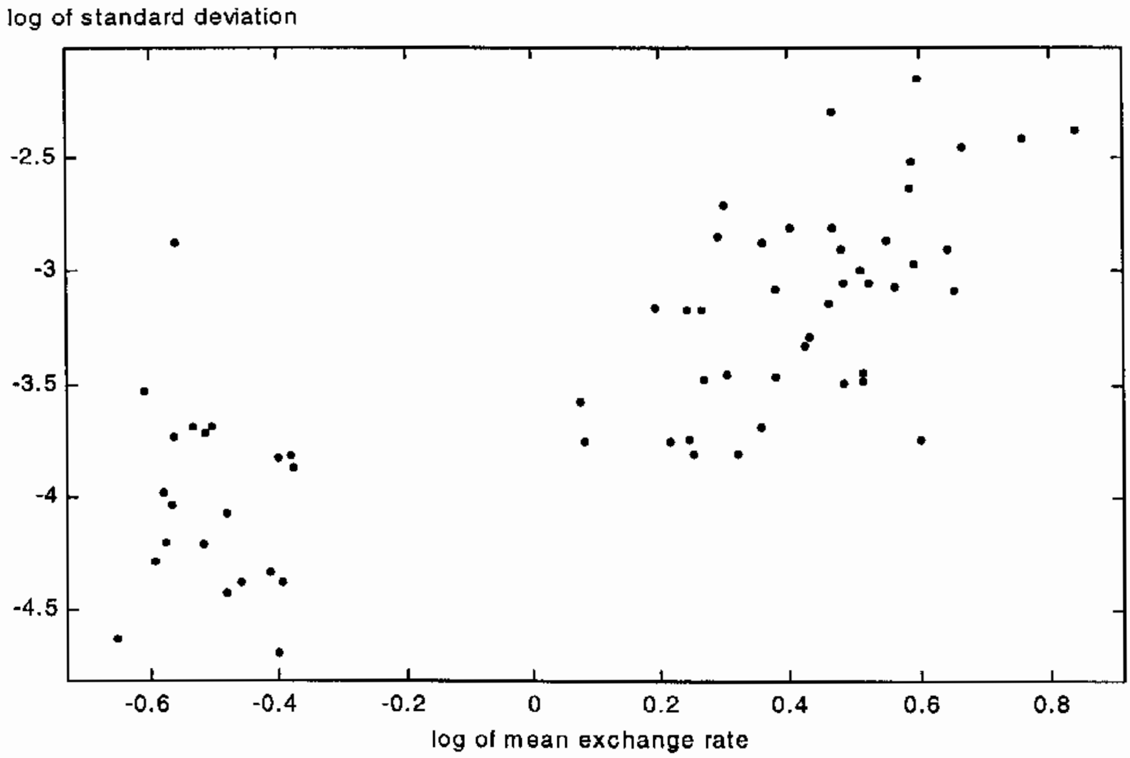

Figure 3.3: Plot of standard deviation versus mean using log transformation

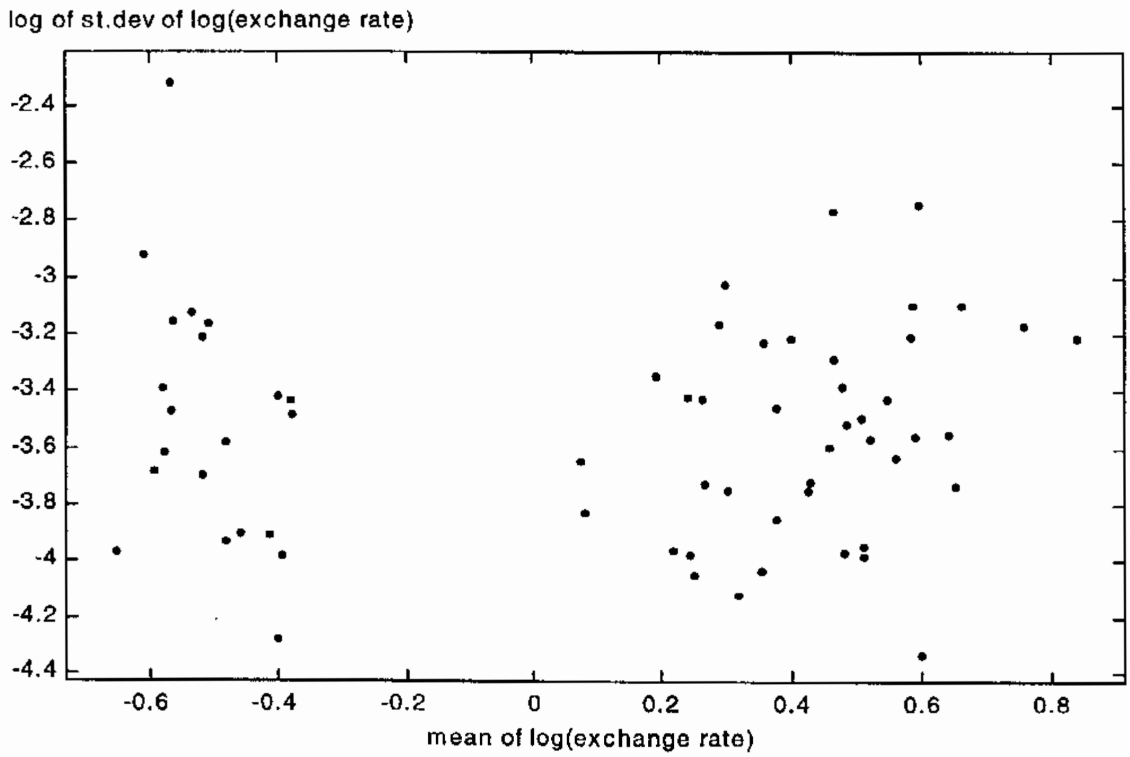

In Figure 3.4 we plot the time series again using the log transformation to these data.

Figure 3.4: Daily exchange rates, transformed by taking natural logarithms

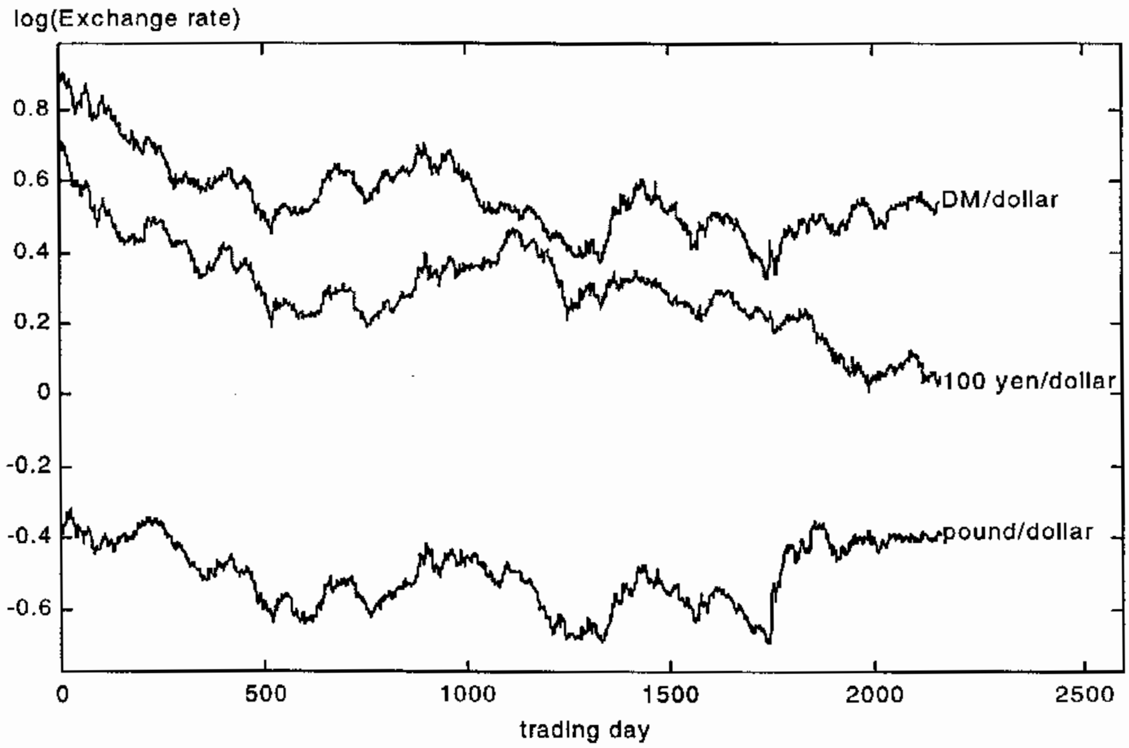

Figure 3.5 highlights the fluctuations in the time series of the pound against the dollar. There is not much between the deutsche mark/dollar and yen/dollar when compared to the raw data time series.

Time series analysis

Next we look at the time series analysis. Figures 3.5 shows the periodogram (both scaled and log-transformed) and the autocorrelation function for the pound against the dollar, for the four-year period from 3 January 1986 to 31 December 1989.

The periodogram shows high energy at low frequencies. The Ljung-Box-Jenkins p-values are all smaller than 0.05 so the fit is not adequate.

Since the log-transformed periodogram has a strong downward trend, we try using a simple autoregressive process to model these data.

Figure 3.6 shows the time series analysis of the residuals, after fitting a first-order autoregressive process.

Figure 3.5: Time series analysis of logarithm of pound/dollar exchange rate, 1986-89

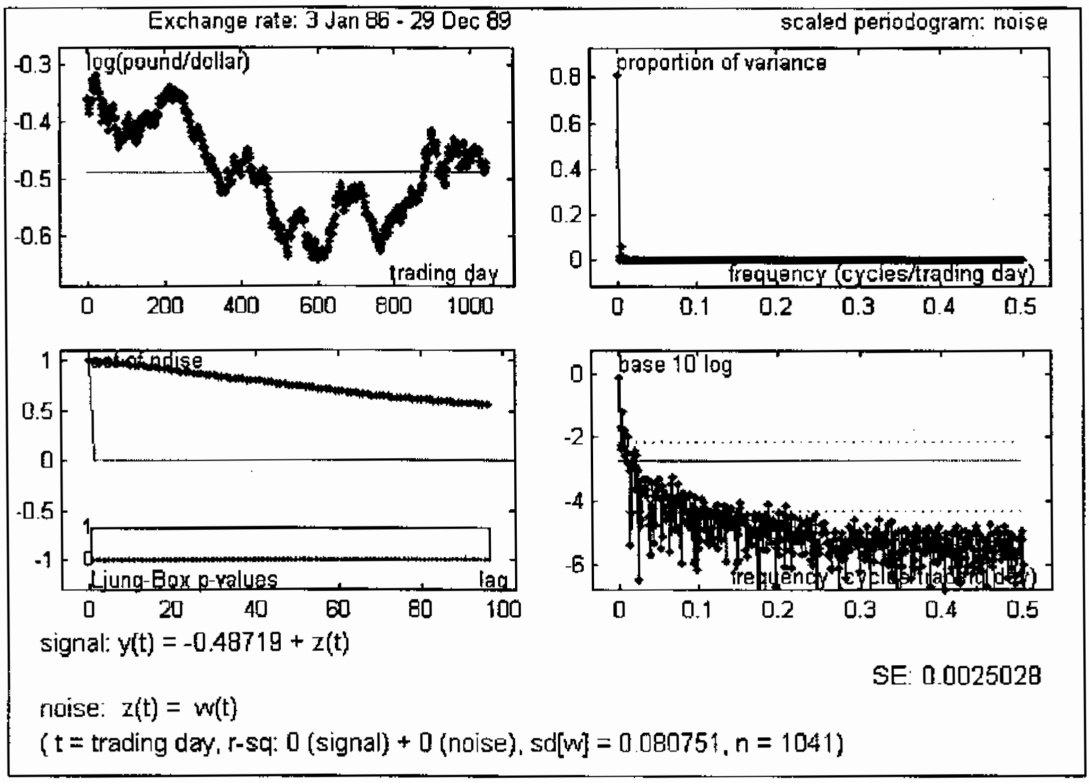

Figure 3.6: Time series of residuals of logarithm of pound/dollar exchange rate, 1986-89, after fitting a first-order autoregression

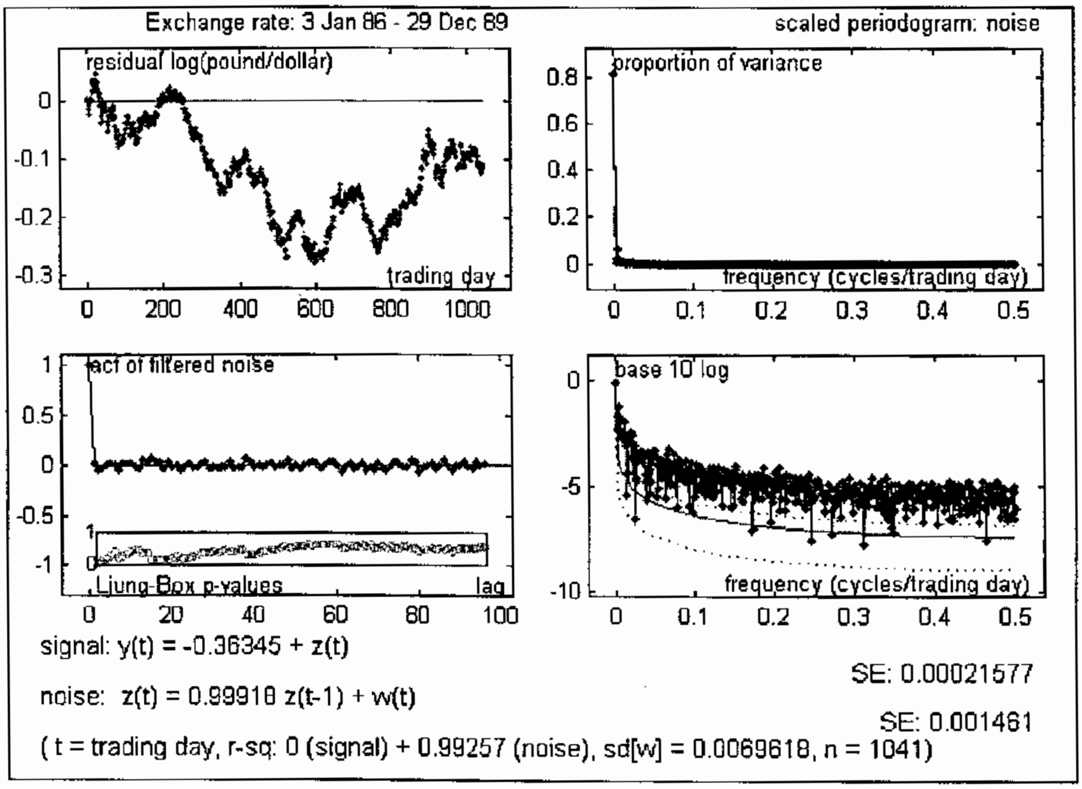

While the periodogram is still outside the confidence limits, all of the Ljung-Box p-values are now greater than 0.05. This indicates that the first-order autoregression provides a good fit. The estimate of the parameter a_1 (0.999) is very close to 1, with standard error 0.001. This suggests that the data should be differenced.

In future, we work with the differences of the logarithms of the exchange rates. These quantities are called *compounded returns.*

Analysis of compounded returns

Table 3.2 shows extended numerical summaries of the daily percentage compounded returns for the three exchange rates compared to the US dollar. These are calculated using Equation (17).

Table 3.2: Numerical summaries of exchange rate returns

col	variable	size	mean	st dev	SE	skew	kurt	min	max
1	pound/dollar	2157	-0.001	0.736	0.016	0.247	5.514	-3.409	4.276
2	100 yen/dollar	2157	-0.031	0.745	0.016	-0.174	5.363	-4.736	3.675
3	DM/dollar	2157	-0.017	0.766	0.016	0.097	5.29	-3.97	3.287

Exchange rates: daily % returns 4 Jan 86 - 12 Apr 94

Table 3.2 shows that the means for the three exchange rates ratios are all negative with similar standard deviations. The kurtoses are all substantially higher that 3. Since the theoretical kurtosis coefficient for a normal distribution is 3, this means that the distributions are longer tailed than the normal distribution.

We now look at the correlations between the exchange rate returns using each of the four currencies in turn as the reference currency. Figure 3.7 shows the scatterplot matrix relating the pound, yen and deutsche mark, using the US dollar as the reference currency. The other correlations, obtained using each of the other currencies as the reference, are shown in Table 3.3.

Figure 3.7 shows that the exchange rates between the three currencies relative to the US dollar are highly positively correlated.

Figure 3.7: Scatterplot matrix showing relations between exchange rate returns relative to US dollar

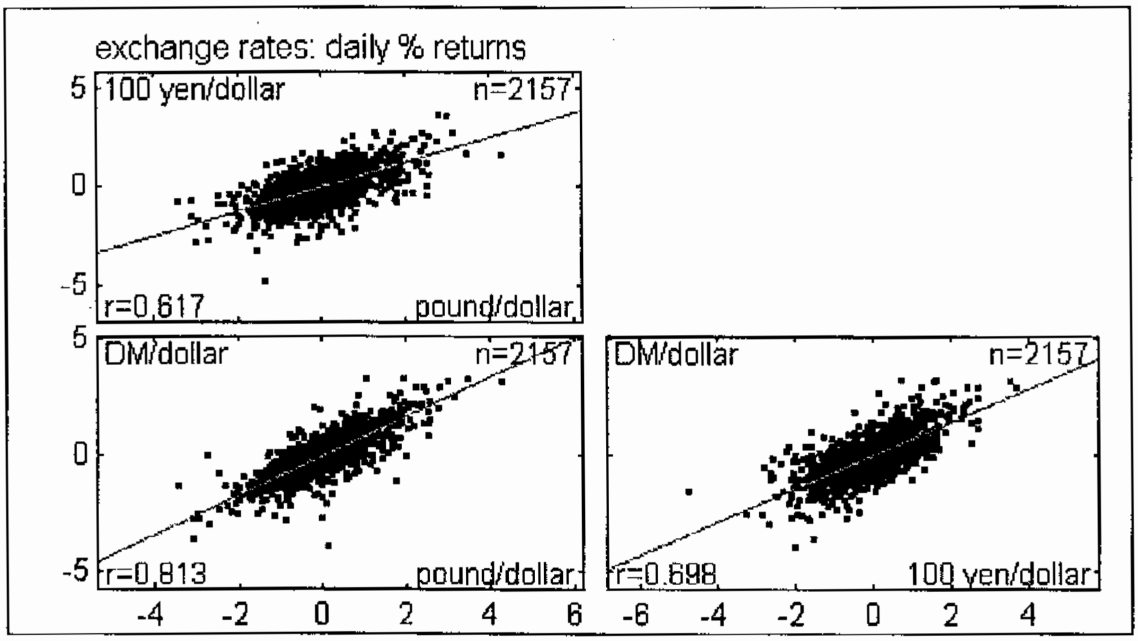

Table 3.3 shows the correlations between all four currency exchange rate returns, taking each currency in turn as the referent currency.

Table 3.3: Correlation between exchange rates (with referent currency in parentheses)

currency	dollar	pound	DM
pound	0.448 (yen) 0.364 (DM)		
DM	0.246 (pound) 0.358 (yen)	0.813 (dollar) 0.727 (yen)	
yen	0.426 (pound) 0.418 (DM)	0.617 (dollar) 0.252 (DM)	0.698 (dollar) 0.481 (pound)

This shows that the correlations are consistent, except for that between the yen and the pound (0.617 or 0.252), using the dollar and deutsche mark, respectively, as the referent exchange rate.

We discuss this further in Chapter 5.

Figure 3.8 shows histograms of the three series of returns relative to the US dollar, with fitted normal distributions superimposed, using the same scale in each case. As seen in Table 3.2, these data range from a minimum of –4.736 (for the yen) to a maximum of 4.276 (for the pound).

The histograms are similar for the three currencies, and clearly show that each distribution is more peaked in the middle than the fitted normal distribution. The kurtoses are greater than 3 in each case, which means that they have longer tails than a normal distribution.

Note that the value for the kurtosis of the deutsche mark is 5.29, in contrast to the obviously incorrect value 36.2 (relative to the pound) given by Shephard (1996, page 3).

Figure 3.8: Histograms of returns

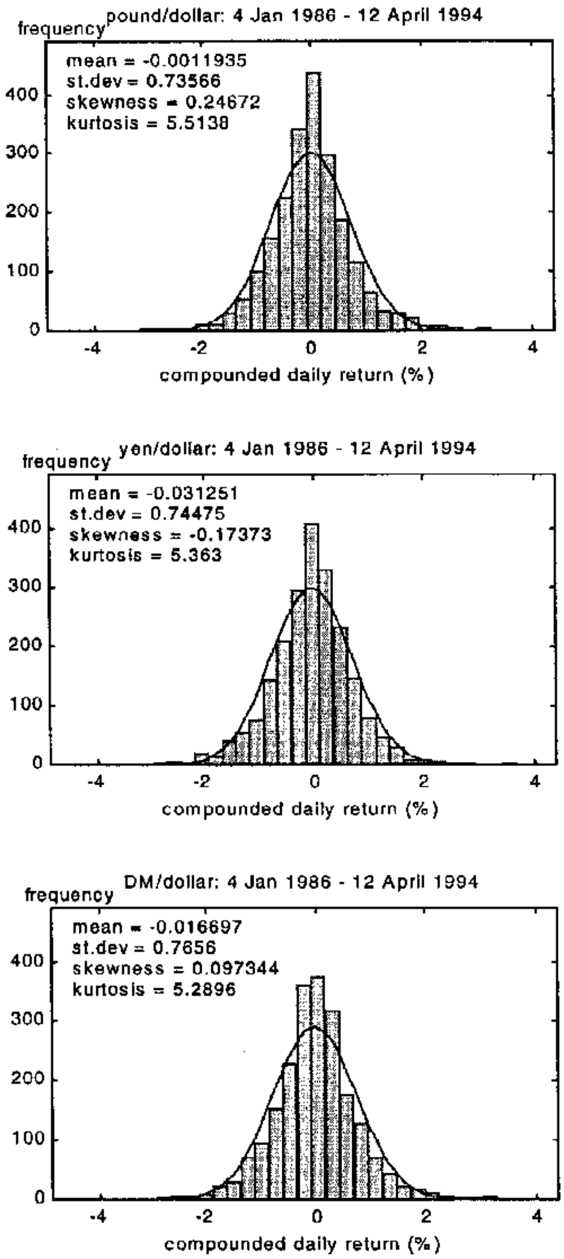

Figure 3.9 shows the ribbon plots (as described in Watcharin thesis) of 95% confidence intervals for exchange rate return means for periods of 50 days, using the US dollar as the referent currency.

These ribbon plots show that the 95% confidence intervals include zero for nearly all of the data. Consequently, the mean is not a useful predictor of the exchange rate.

Figure 3.9: Ribbon plots of 95% confidence intervals for exchange rate returns means for 50-day periods

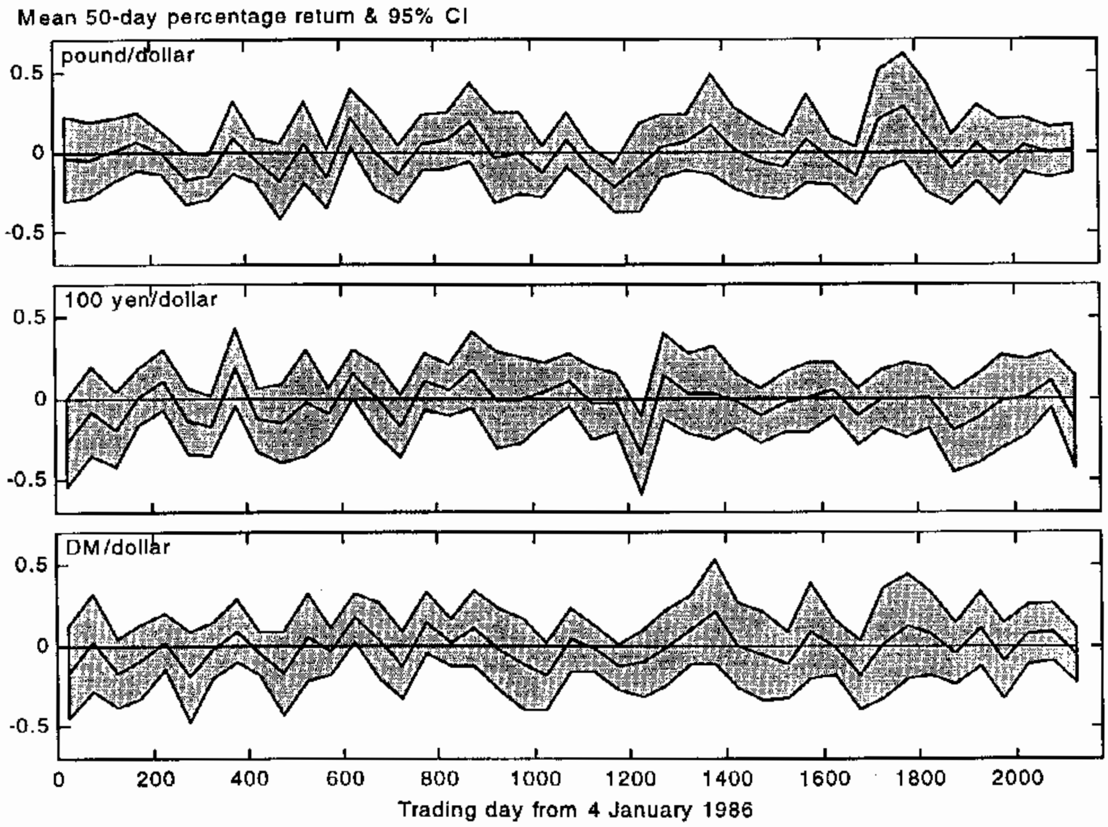

Next we plot the volatility (estimated standard deviation) of these data, as shown in Figure 3.10, using periods of 25 days and 50 days to estimate the volatility.

Figure 3.10 shows that the volatility of each currency tends to follow the same pattern, with overall standard deviations that are very close together (0.736 for the pound, 0.745 for the yen and 0.766 for the deutsche mark). Clearly, the pound and the deutsche mark show similar high fluctuations in the same period. We see the series moving together, indicating that the exchange rate returns are self-correlated.

Similarly, Figure 3.11 shows that the estimated skewness coefficient of each currency exchange rate return tends follow the same pattern. However, the overall skewness coefficients differ for the three exchange rates, having positive values for the pound (0.25) and the deutsche mark (0.10), and a negative value for the yen (-0.17). For each exchange rate, the estimated skewness in the returns dropped substantially just before the 600th trading day, corresponding to the stock market crash of 19 October 1987.

Figure 3.10: The volatility of standard deviations for exchange rate returns

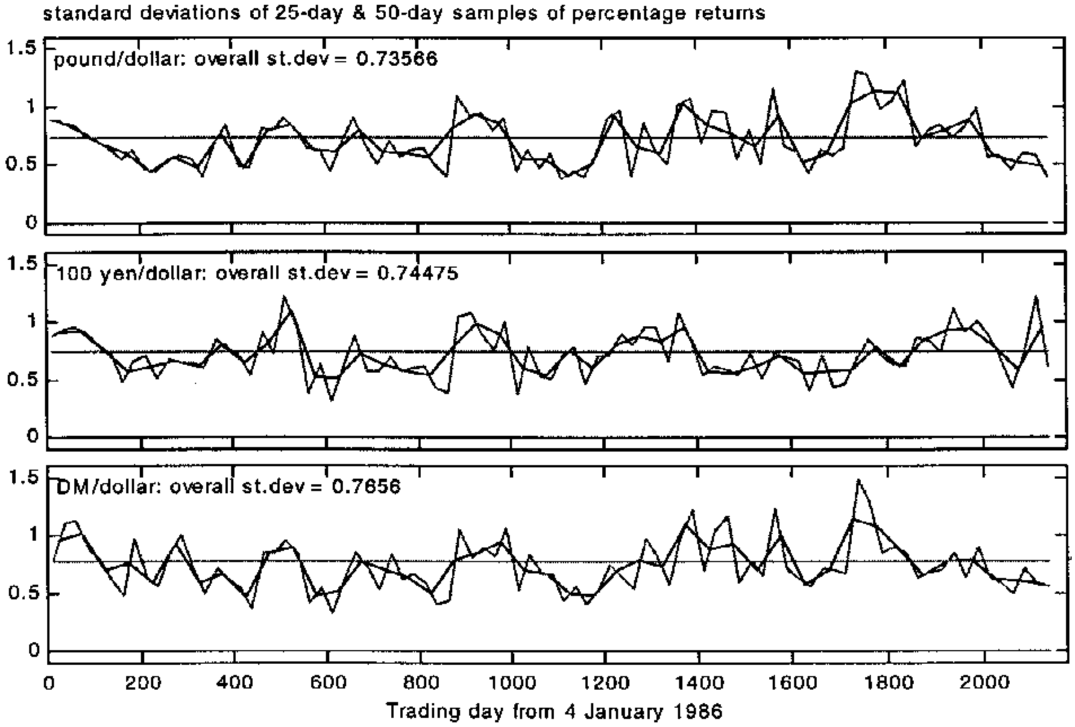

Figure 3.11: The volatility of skewnesses for exchange rate returns

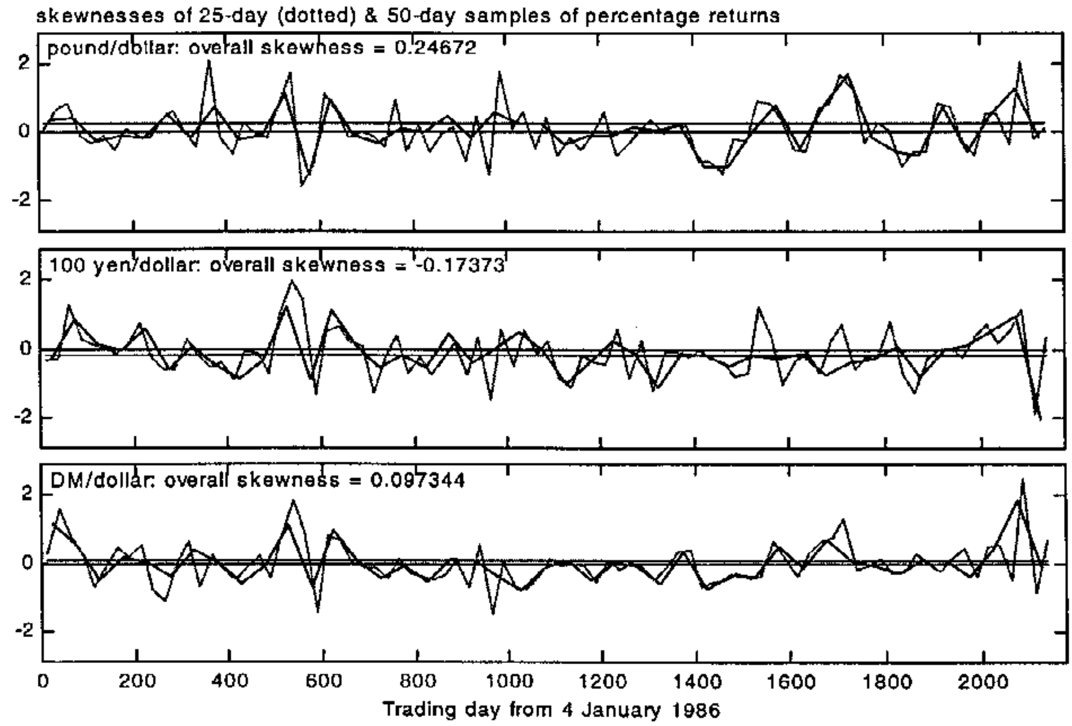

Figure 3.12 shows that the volatility kurtosis of each currency exchange rate return tends to follow the same pattern, with overall kurtosis coefficients that are very close together. These overall values are 5.51 for the pound, 5.36 for the yen, and 5.29 for the deutsche mark. Not surprisingly, the crash of 19 October 1987 is marked by a substantial increase in volatility. However, there are even higher increases in kurtosis in each series (particularly for the yen) near the end of the study period.

Figure 3.12: The volatility of kurtoses for exchange rate returns

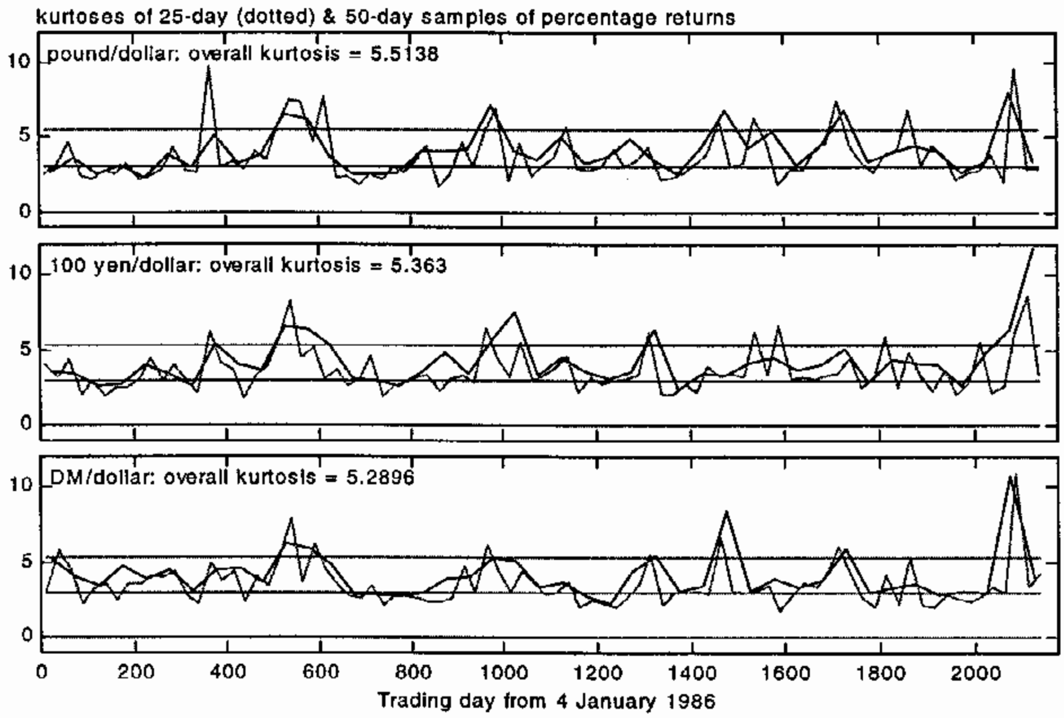

Figure 3.13 shows the relations between the volatilities in the three series of exchange rate returns, depicted as a scatterplot matrix. The correlations are all positive. The highest correlation occurs between deutsche mark and the pound with the value 0.779. The correlations for other two relations are similar but lower, with value 0.397 and 0.389, respectively, for the yen versus the pound and the yen versus the deutsche mark.

Figure 3.14 shows that the correlations of the skewness coefficients for the exchange rate returns are again positively correlated. The highest correlation again occurs between the deutsche mark and the pound with the value 0.599. And as for the volatilities, the correlations for the other two relations are similar, with values 0.392 and 0.454.

Figure 3.13: The scatterplot matrix of volatilities for exchange rate returns

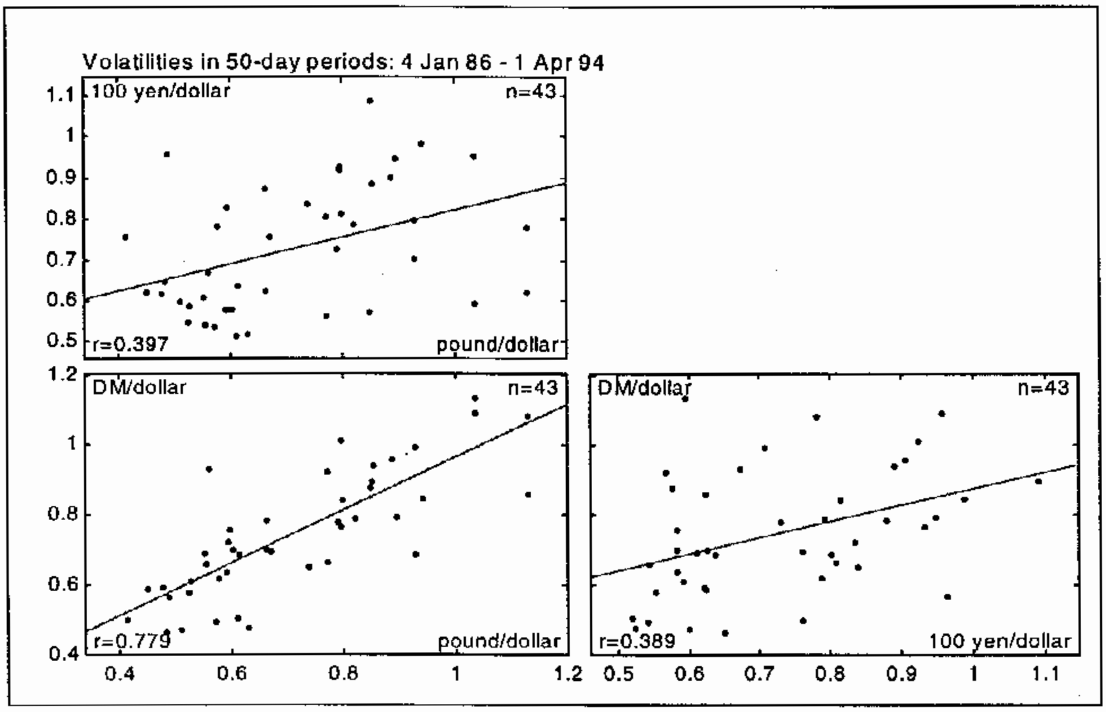

Figure 3.14: The scatterplot matrix of the skewnesses for exchange rate returns

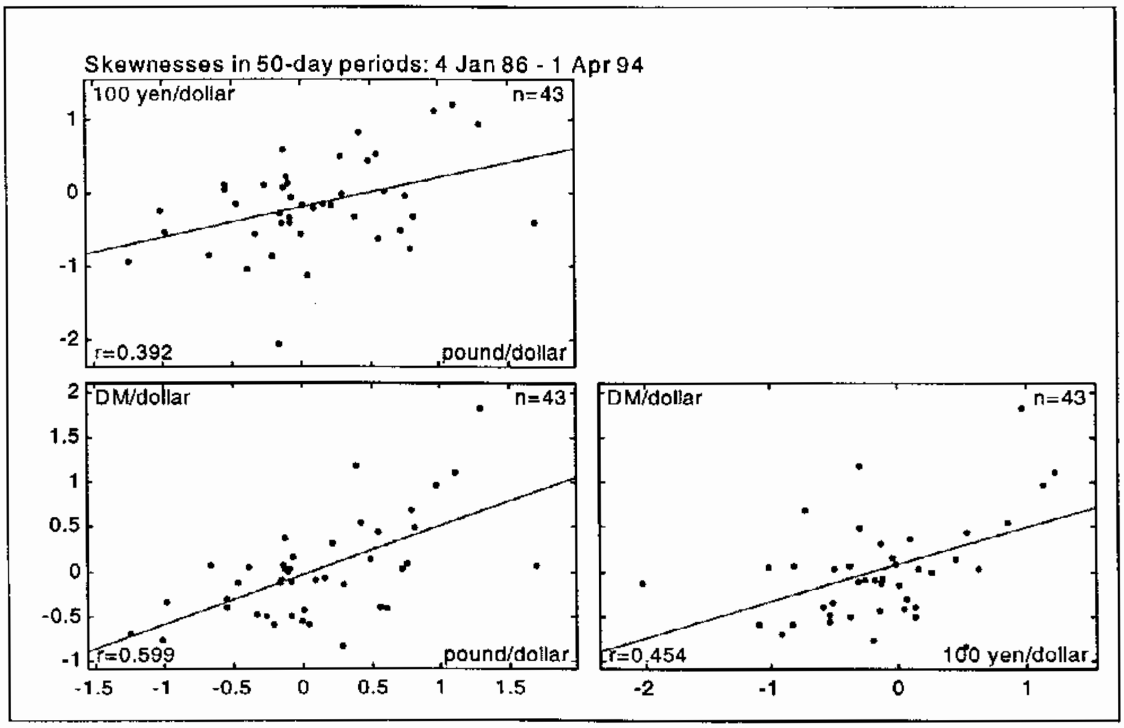

Figure 3.15 shows the scatterplot matrix and the correlation values for the kurtosis coefficents, before and after using a logarithm transformation for these estimates.

Figure 3.15: Scatterplot matrix of the kurtosis coefficients for exchange rate returns

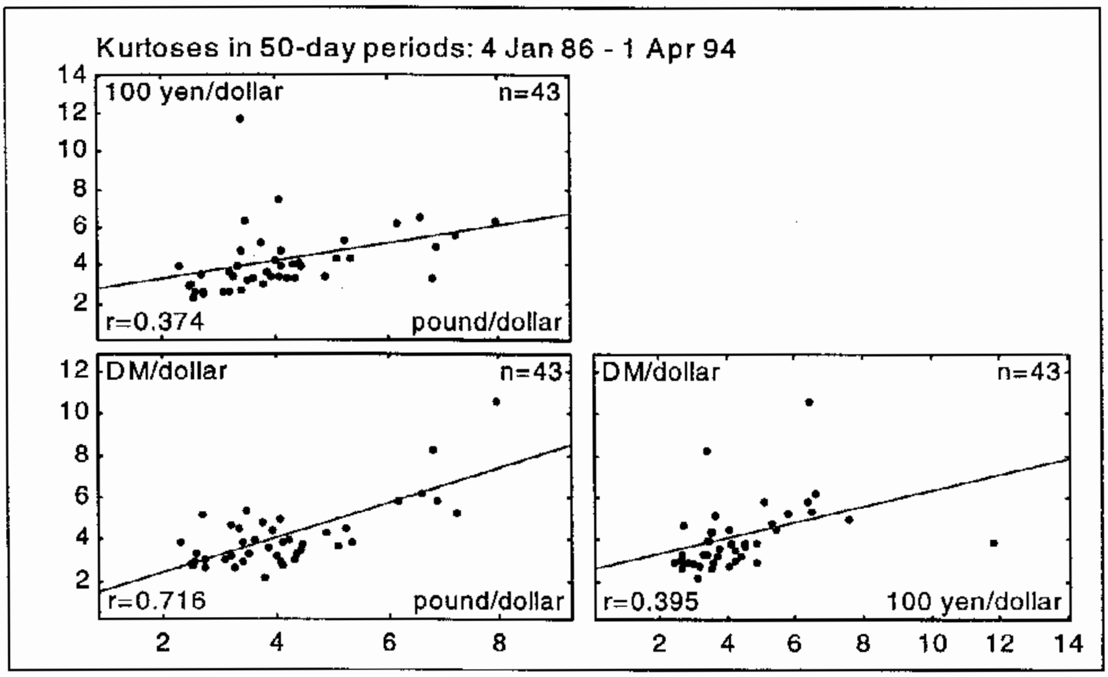

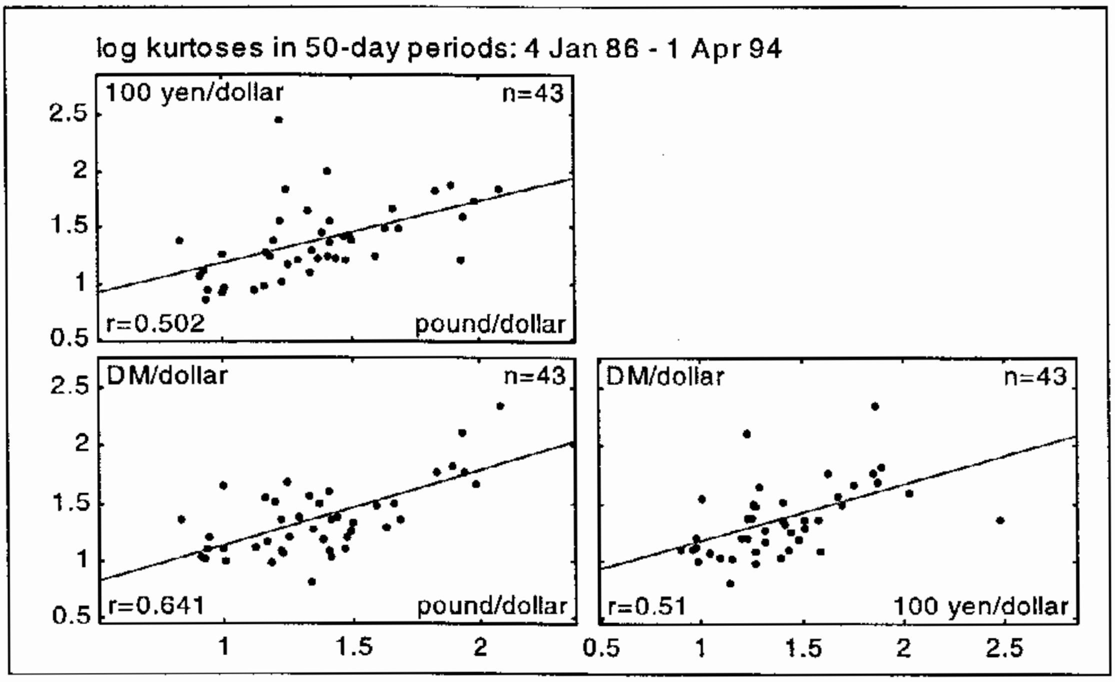

Figure 3.15 shows that the correlations of the kurtosis volatility (top), and after taking logarithms of the kurtosis (bottom) are positively correlated to each currency exchange rate. As for the other coefficients, both graphs show the highest correlation occurs between the deutsche mark and the pound.

Summary

The preliminary results show that the logarithm transformation is appropriate for analysing the time series of exchange rates for the pound, yen and deutsche mark by using the US dollar as the referent.

First we look at the time series analysis, which suggests differencing these data. These differences are called compounded returns.

Second we analyze the compounded returns. These show that there are positively correlation between the four currency exchange rates, by using four of them in turn as the reference. Most of the correlations are consistent. The exceptional correlations are between the pound and the yen, using the US dollar and deutsche mark as referents.

Third we plot the histograms of these exchange rates. The histograms are similar for the three currencies, and clearly show that each distribution is more peaked in the middle than the fitted normal distribution. The kurtoses are greater than 3 in each case, which means that they have longer tails than a normal distribution.

Fourth we plot the ribbon plot of means. It shows that mean is not a predictor for exchange rate. Then we plot the volatility (or standard deviation), skewness and kurtosis for these three currencies. The volatility of each currency tends follow the same pattern. However, the pound and the deutsche mark move together.

Finally we look at the correlations in the standard deviation, skewness and kurtosis. These show evidence of high correlations between the pound and the deutsche mark.