Chapter 4

Estimation and Simulation

In Chapter 3 we examined the time series of daily exchange rate returns for the pound, the yen, and the deutsche mark (relative to the US dollar), over the period from the beginning of January 1986 until April 1994. We found the three series were correlated with each other, and their distributions were longer tailed than the normal distribution. By examining the moments of the distributions within 25-day and 50-day stretches, we found that the three series of exchange rate returns were similar. We also found evidence of non-constant volatility.

In this chapter we estimate the parameters in the simple model for stochastic volatility fitted to the data. We then assess the model by creating simulated time series with these parameters, and thus compare the simulated series with the data.

Time series analysis of volatility

Refer to Figure 3.10, which shows the changing volatility of the three exchange rate returns, based on 25-day and 50-day averages. First, we test the null hypothesis that the volatility is constant.

Figure 4.1 gives a time series analysis of the volatilities, using periods of 25 days, for the pound/dollar exchange rate. This series is fitted well by a first-order autoregression with parameter $a = 0.368$. The standard error of this estimate is 0.102, so the z-statistic for testing the null hypothesis is $z = 0.368/0.102 = 3.61$, giving a small p-value (0.0003).

Now if the volatility were constant, the time series of volatility estimates would comprise a sequence of independent quantities, and would thus resemble white noise. Since the volatility series is fitted by a first-order autoregression with parameter greater than 0, we conclude that the volatility of the pound/dollar exchange rate return is not constant.

Figures 4.2 and 4.3 show the corresponding time series analyses for the yen/dollar and deutsche mark/dollar returns, again using periods of 25 days to estimate the volatilities.

Figure 4.1: Time series analysis of 25-day volatility for pound/dollar return

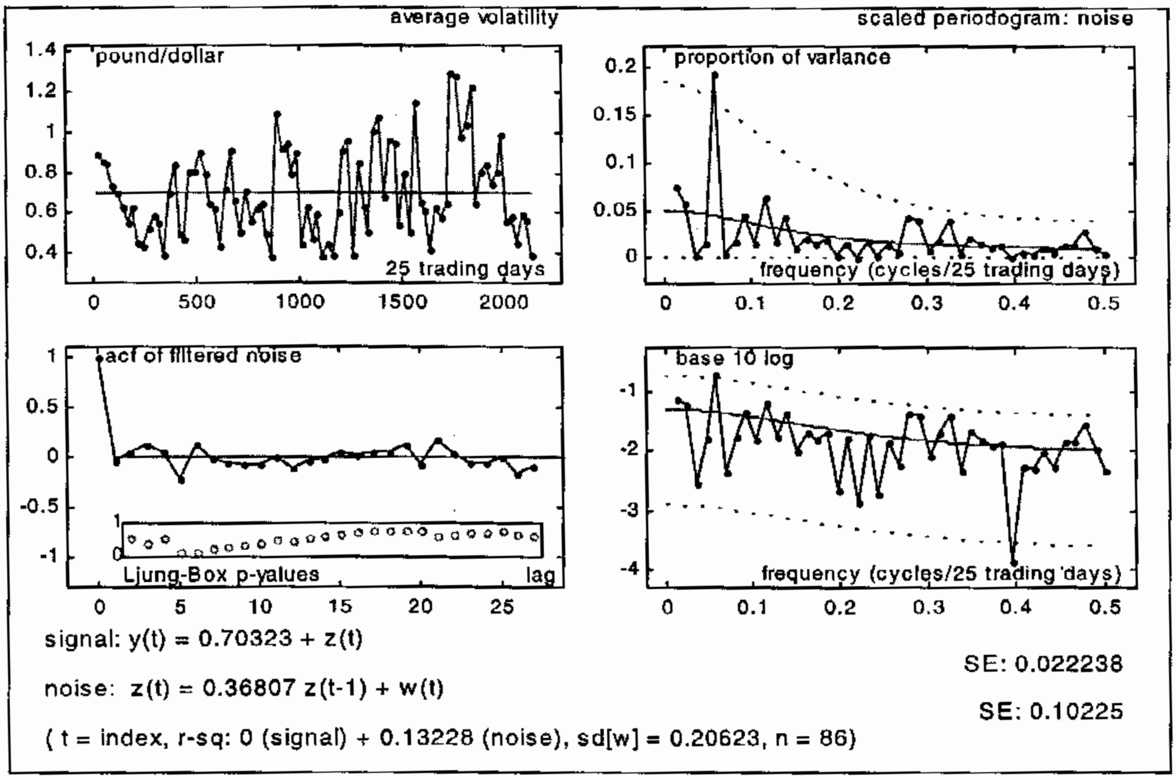

Figure 4.2: Time series analysis of 25-day volatility for yen/dollar return

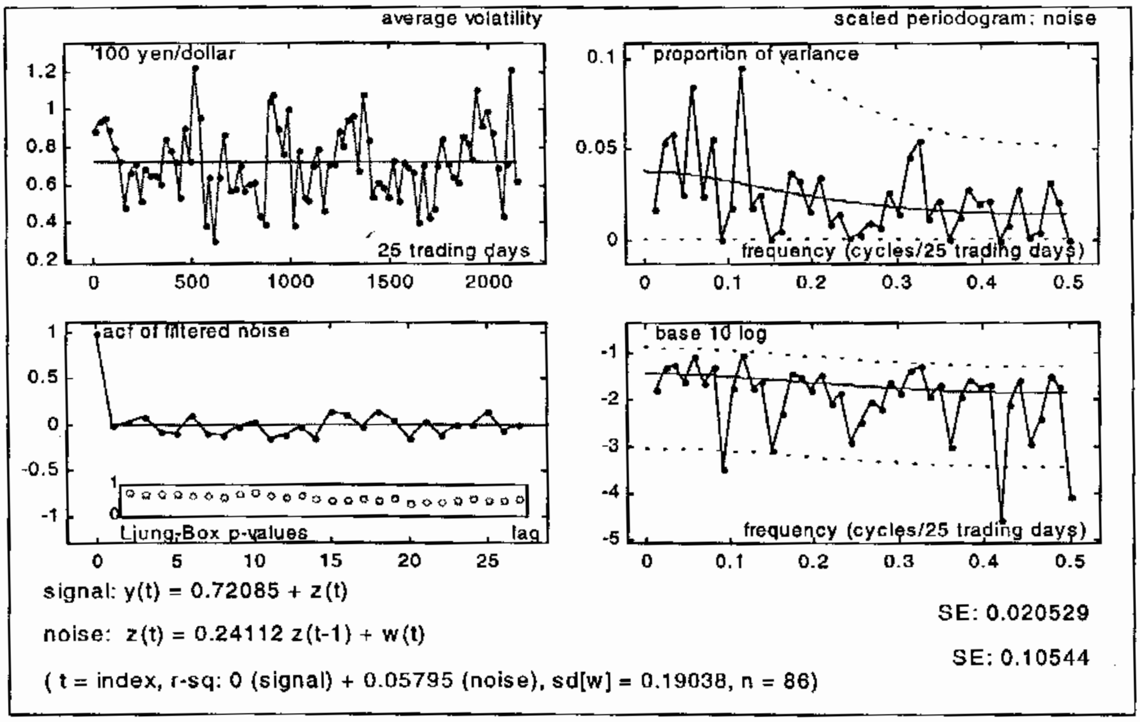

Figure 4.3: Time series analysis of 25-day volatility for DM/dollar return

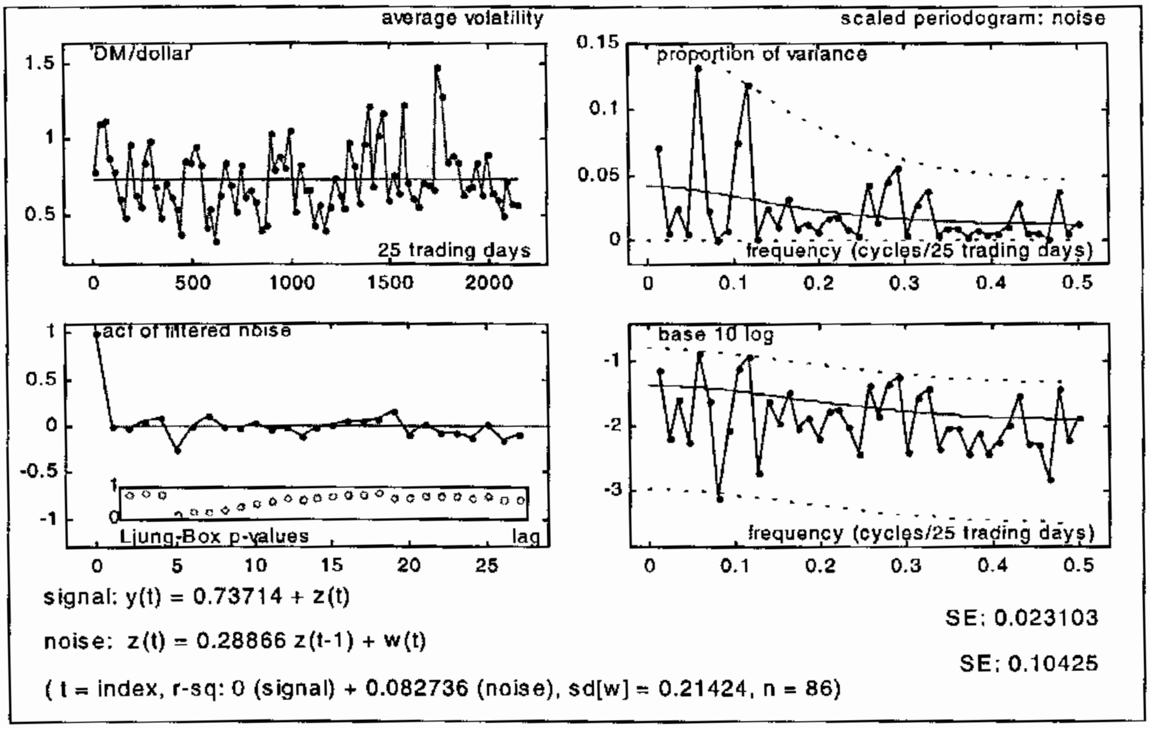

For these currencies, the volatility series is again fitted well in each case by a first-order autoregression, with parameters 0.241 and 0.289, respectively. Since the standard errors of these estimates are 0.105 and 0.104, respectively, the z-statistics for testing the null hypothesis of constant volatility are 2.30 and 2.78, respectively. The corresponding p-values are 0.021 and 0.0054, so the null hypothesis is rejected in each case.

We have used periods of 25 days to estimate the volatility of each series. Table 4.1 shows results for periods ranging from 20 to 40 days. The highest autocorrelation coefficient is obtained in each case when a period of 25 trading days is used to estimate the volatility.

Table 4.1: Estimates and standard errors of autoregressive parameters

Period	Pound/dollar		100 yen/dollar		DM/dollar	
	autocorrelation	SE	Autocorrelation	SE	autocorrelation	SE
20	0.322	0.093	0.186	0.095	0.269	0.094
25	0.368	0.102	0.241	0.105	0.289	0.104
30	0.313	0.114	0.233	0.121	0.158	0.119
40	0.268	0.134	0.045	0.144	0.144	0.137

Estimation of parameters

We have shown that the volatility of the exchange rate returns is not constant. So we now try to fit the stochastic model developed in Chapter 2, given by equations (20) and (21). The moments of the distribution of a process that follows this model are given by equations (26) – (29), and the autocorrelations of the squares of the time series at different lags are given by equation (30).

Since these equations are nonlinear, a numerical inversion procedure is required in general. However, their solution is much simpler if we assume that the parameter α is close to 1, so we can use equations (31) – (35).

For the three series of exchange rate returns (for the pound, yen, and deutsche mark relative to the dollar), the estimated standard deviations for the period from 4 January 1986 to 12 April 1994 were found to be 0.736, 0.745 and 0.766, respectively (see Figure 3.12). The average is thus 0.749. The corresponding estimates of the overall kurtosis over this period were 5.514, 5.363, and 5.290, respectively, averaging to 5.389. Since the overall means are statistically indistinguishable from 0, we assume that $\mu = 0$. Substituting these average values into equations (36) – (38), we find

$$\hat{\sigma}^2 = 0.749^2 \sqrt{(1.5 - 5.389/6)} = 0.435,$$

$$\hat{\kappa} = \sqrt{(0.749^2 - \hat{\sigma}^2)} = 0.355.$$

These estimates are the limiting values as $\alpha \to 1$, and at this limit, δ is 0. But since δ must be positive for the process to have variable volatility, so we need to choose a value of α slightly below 1 to get a proper estimate for δ. Since $\delta\beta = \kappa$, where $\beta = \alpha/\sqrt{(1-\alpha^2)}$, it follows that $\delta = \kappa\sqrt{(1-\alpha^2)}/\alpha$.

For example, if $\alpha = 0.9$, $\delta = 0.355\sqrt{(1-0.9^2)}/0.9 = 0.172$. Similarly, if $\alpha = 0.95$, we find that $\delta = 0.355\sqrt{(1-0.95^2)}/0.95 = 0.117$.

The parameter α can be estimated from the autocorrelation function of the squared returns, by fitting the theoretical function given by equation (35).

Figure 4.4 shows this autocorrelation function for the pound/dollar returns, together with the theoretical functions obtained by taking $\alpha = 0.95$ (solid curve) and $\alpha = 0.90$ (dotted curve).

Figure 4.4: Autocorrelation function of squared exchange rate returns for pound/dollar

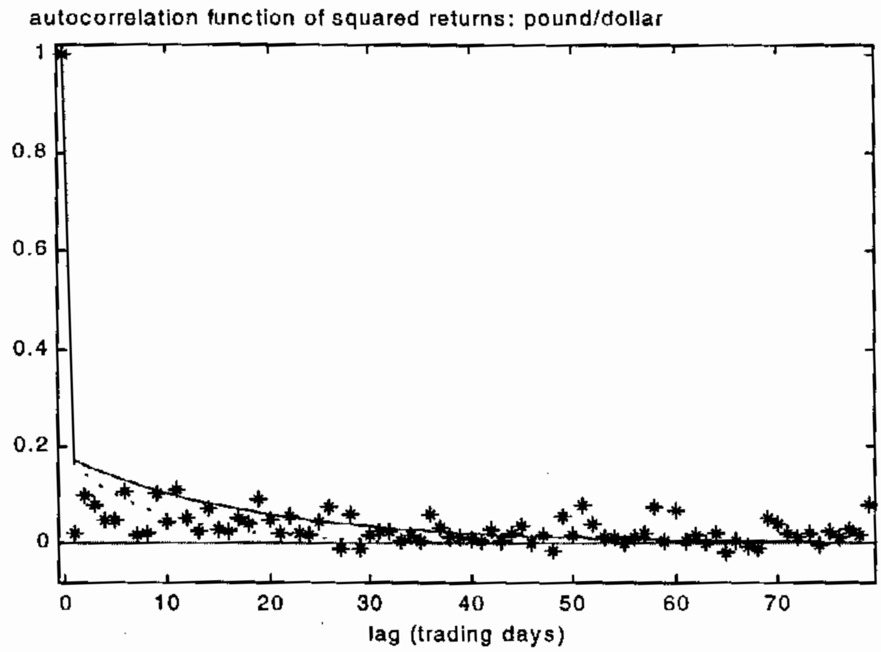

Note that the autocorrelations based on the data are mostly positive. However, for small lags, the theoretical values are higher than the values obtained for the data, particularly for $\alpha = 0.95$. This indicates that the limiting model does not provide a satisfactory fit.

Figure 4.5 shows simulations of the volatility based on 25-day and 50-day stretches, with $\alpha = 0.9$ (top three plots) and $\alpha = 0.95$ (bottom three plots). This figure shows that the estimated standard deviation coefficients of these simulated exchange rate returns (average of 0.7488 for $\alpha = 0.9$, average of 0.7301 for $\alpha = 0.95$) are close to those for the data (average of 0.749).

Figure 4.5: Simulation volatility of standard deviations for exchange rate returns

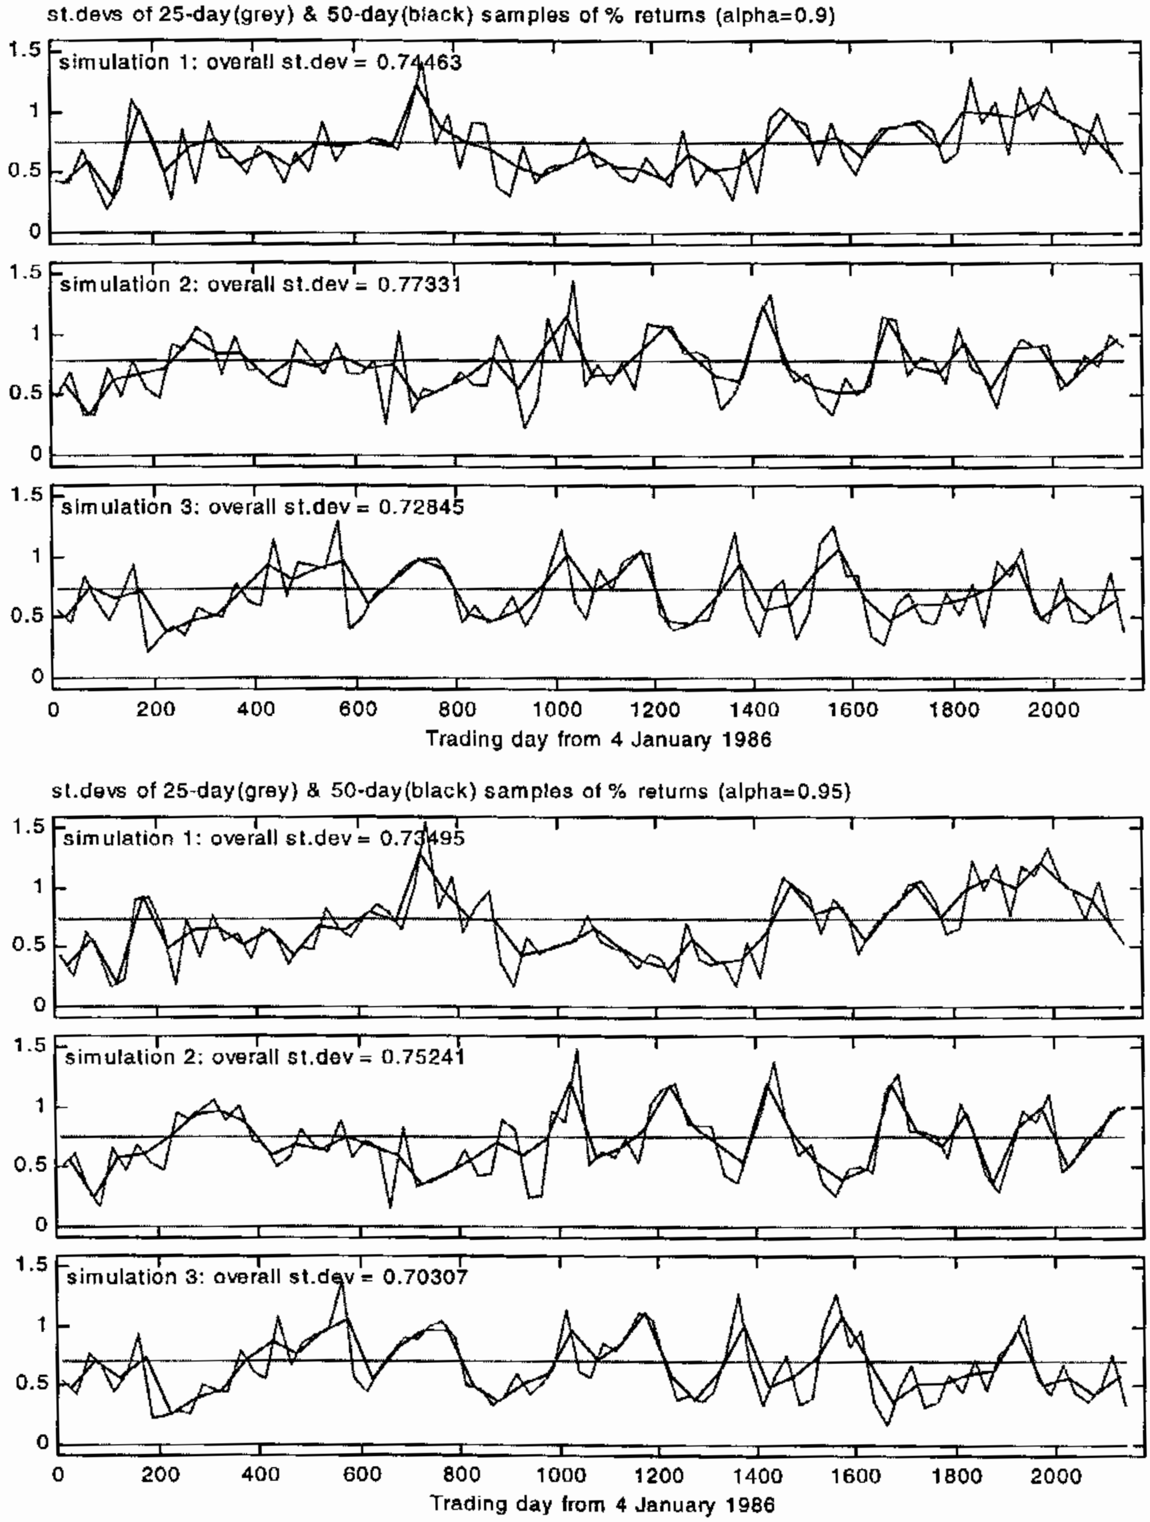

When compared with the standard deviations coefficients of 25 days with $\alpha = 0.9$ shown in Figure 3.10, it can be seen that the volatility series based on the simulation looks just like the volatility series based on the data (Figure 4.6)

Figure 4.6 : Simulation volatility of standard deviations compared with the volatility on the data basis for exchange rate returns

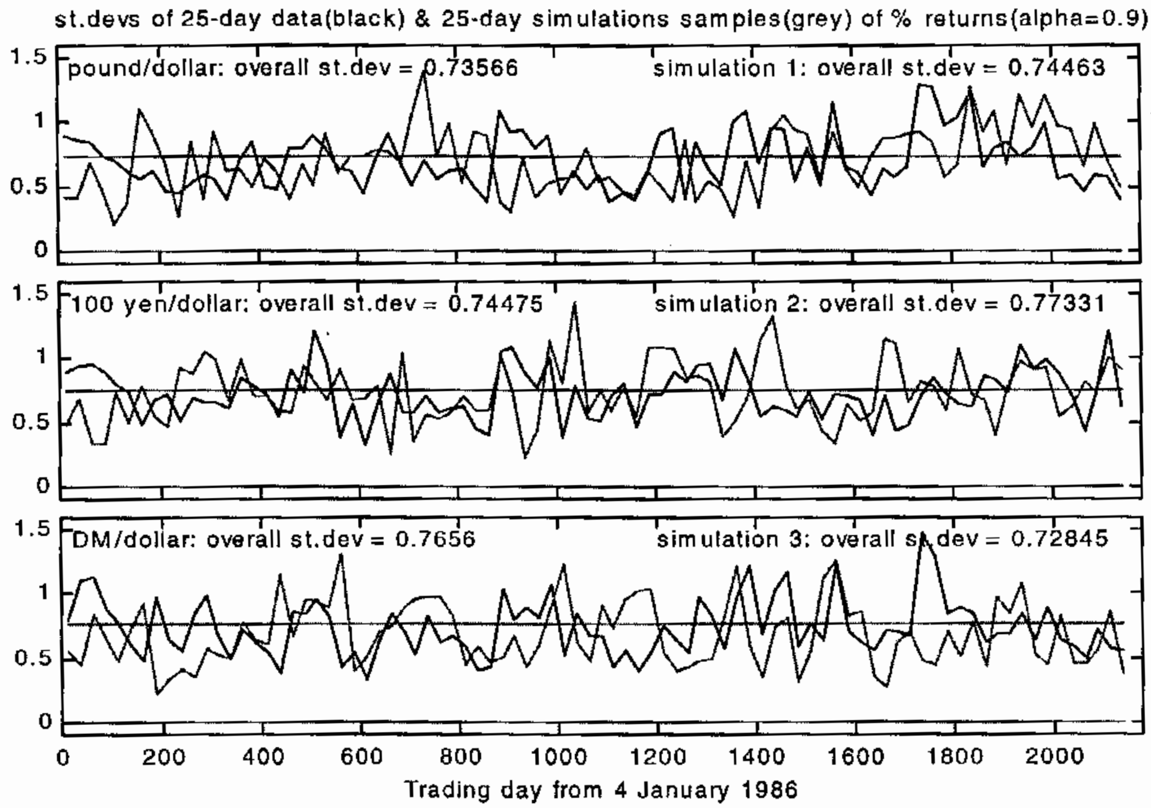

The series of skewnesses based on the simulated data, again using periods of 25 days and 50 days, for $\alpha = 0.9$ (top three plots) and $\alpha = 0.95$ (bottom three plots) (Figure 4.7).

Figure 4.7 shows that the skewness coefficients based on the simulations (range of -0.11 to 0.03, average absolute value of 0.06) are less than the estimate overall skewness coefficients of the data (range of -0.17 to 0.25, average absolute value of 0.17) (see Figure 3.11).

Figure 4.7: Simulation volatility of skewnesses for exchange rate returns

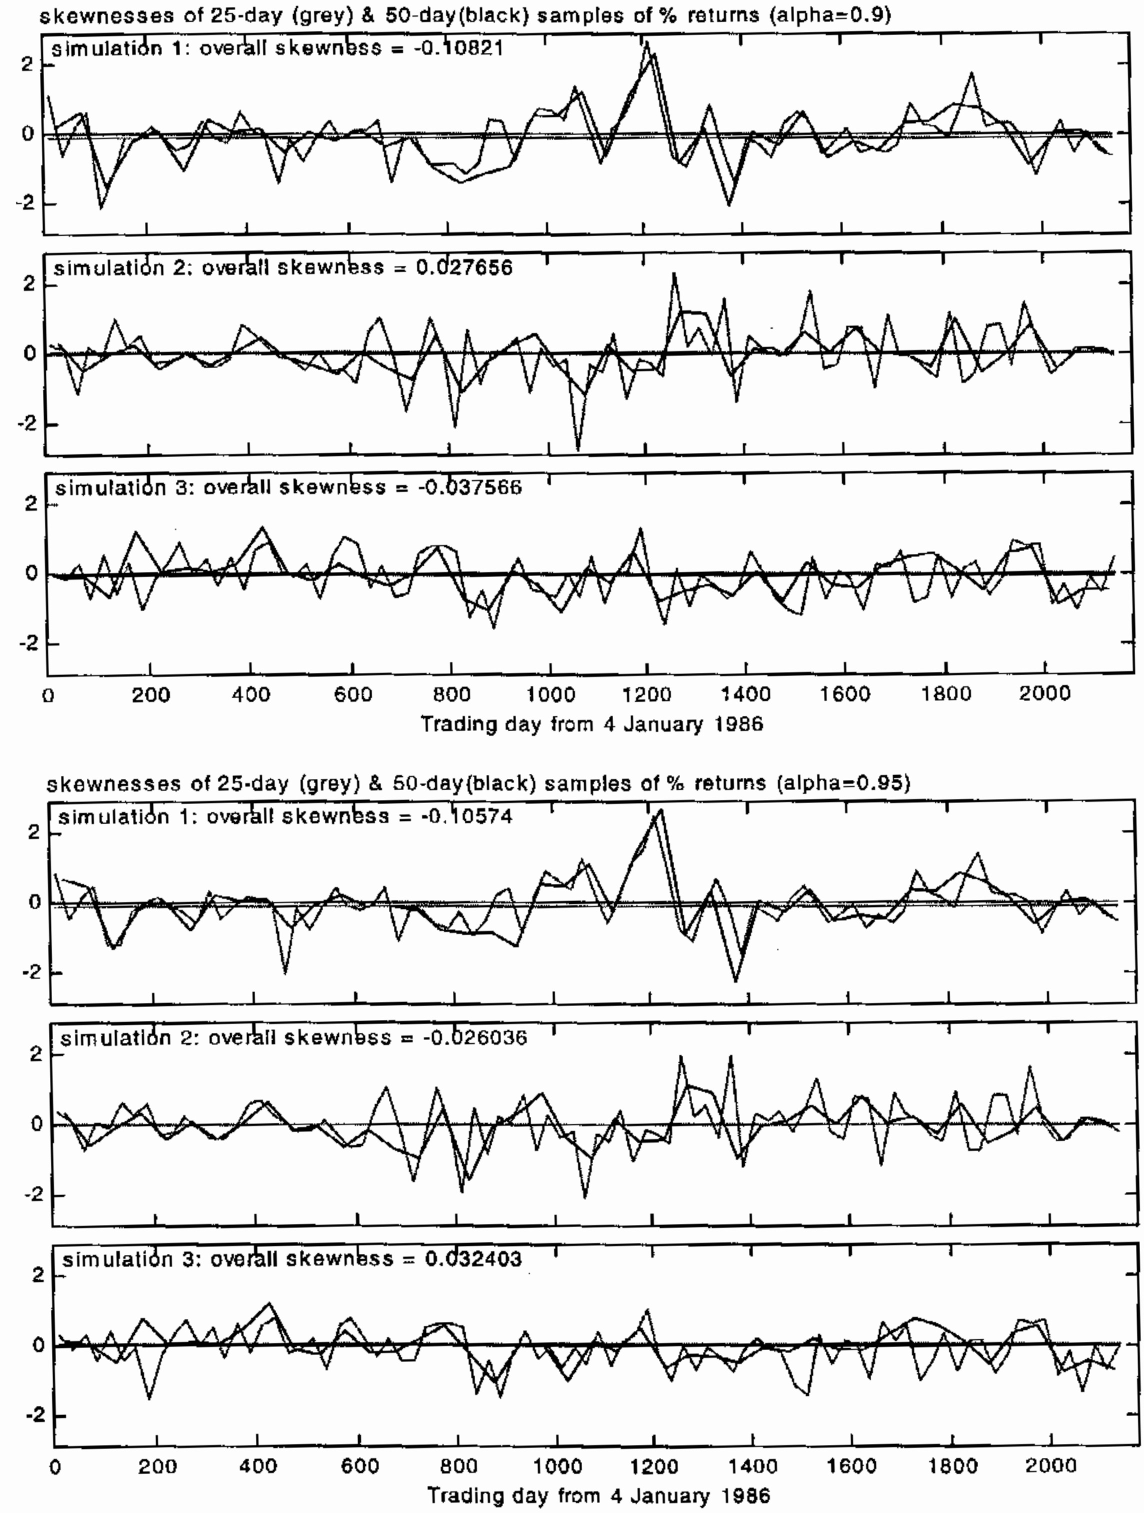

These plots with 25 days and α of 0.9 are compared with those shown in Figure 3.11 as presented in Figure 4.8.

Figure 4.8 : Simulation volatility of skewnesses compared with the volatility on the data basis for exchange rate returns

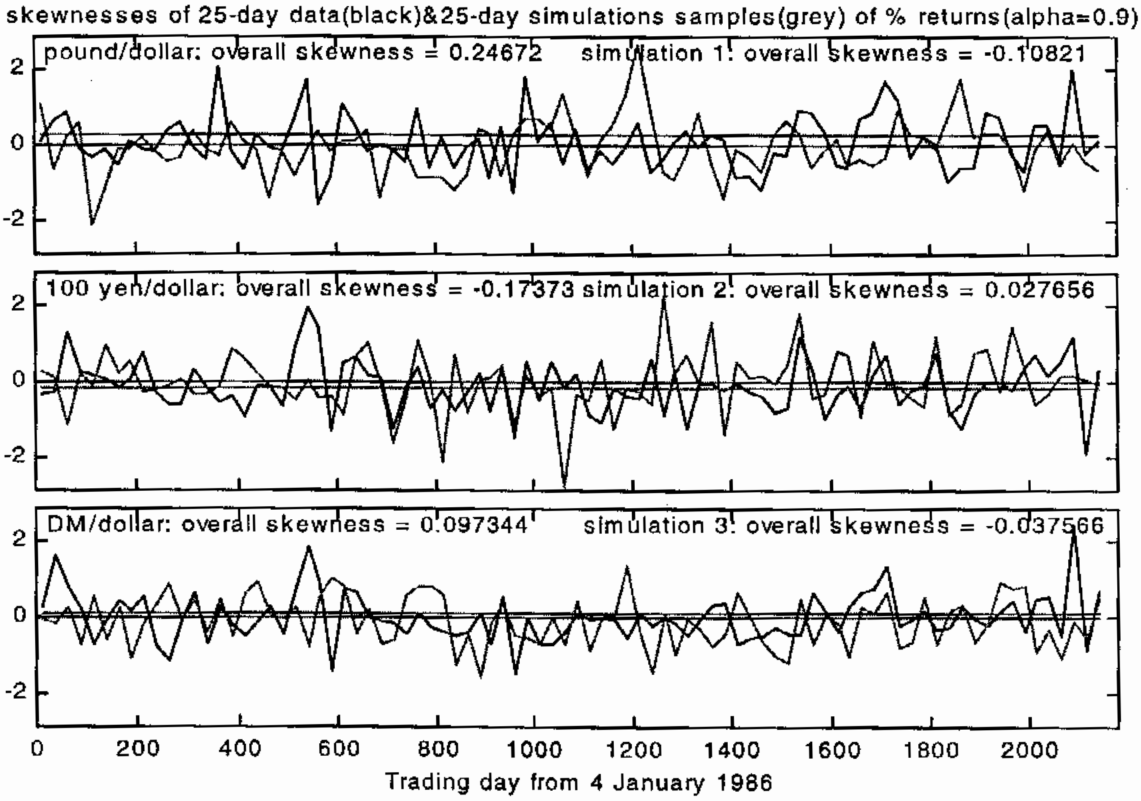

Figure 4.9 show the series of kurtosis coefficients based on the simulations, again using periods of 25 day and 50 days, for $\alpha = 0.9$ (top three plots) and $\alpha = 0.95$ (bottom three plots). These plots should be compared with Figure 3.12. This figure shows that the overall kurtosis coefficients based on the simulations (average 5.38 for $\alpha = 0.9$, 5.40 for $\alpha = 0.95$) are close to the overall kurtosis coefficients of the data (average of 5.39) (Figure 3.12).

Figure 4.9: Simulation volatility of kurtoses for exchange rate returns

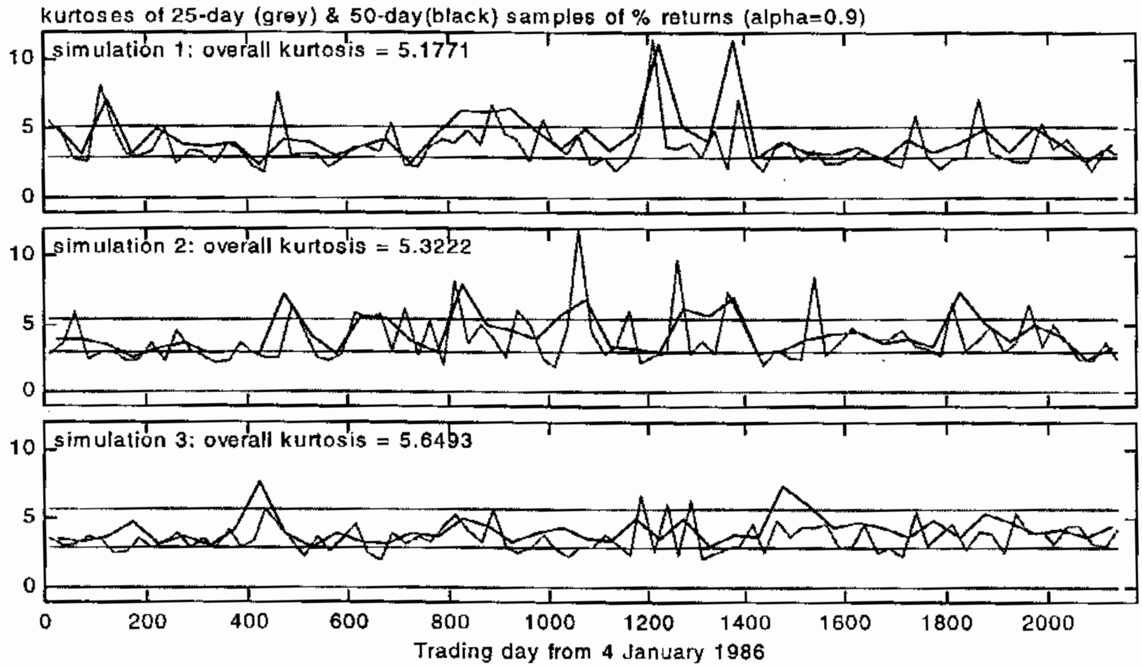

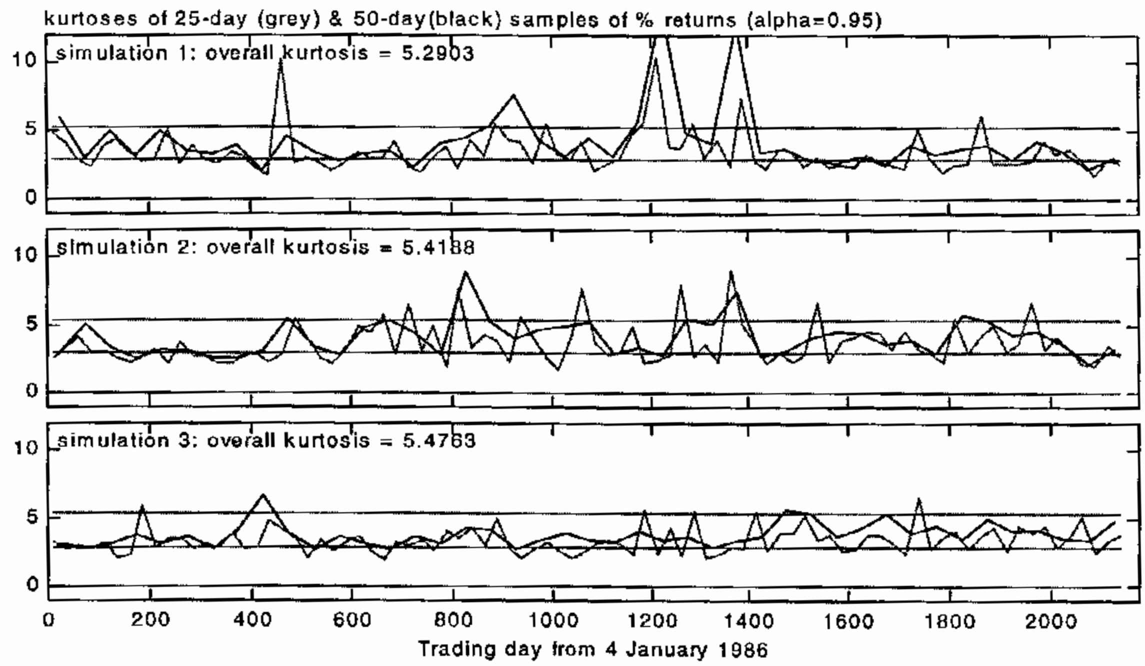

When compared with the kurtosis coefficients of 25-day with $\alpha = 0.9$ shown in Figure 3.12, it can be seen that the kurtosis series based on the simulations is somewhat similar to the kurtosis series based on the data (Figure 4.10).

Figure 4.10 : Simulation volatility of kurtoses compared with the volatility on the data basis for exchange rate returns

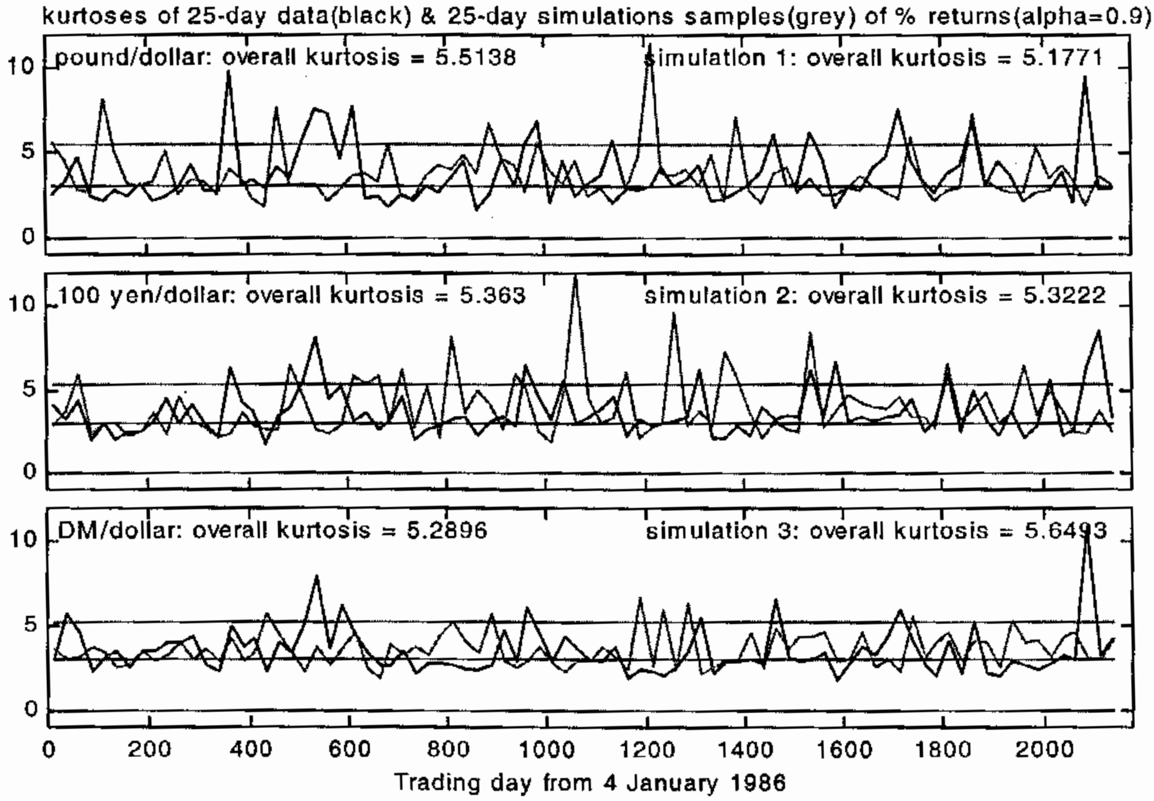

Time series analysis of the simulation volatility

Figures 4.11 and 4.12 show the time series analyses for the volitility series based on the simulated data, for $\alpha = 0.9$ and $\alpha = 0.95$, respectively.

Figure 4.11 shows the time series analysis of the volatility of the first simulation series for $\alpha = 0.9$, for exchange rate returns based on 25-day averages. First, we test the null hypothesis that the simulation of volatility is constant.

This series is fitted well by a first-order autoregression with estimated parameter $a = 0.295$. The standard error of this estimate is 0.104, so the z-statistic for testing the null hypothesis is z = 0.295/0.104 = 2.8365, giving a very small p-value (0.0046).

Figure 4.11: Time series analysis of 25-day volatility for simulated data with $\alpha = 0.9$

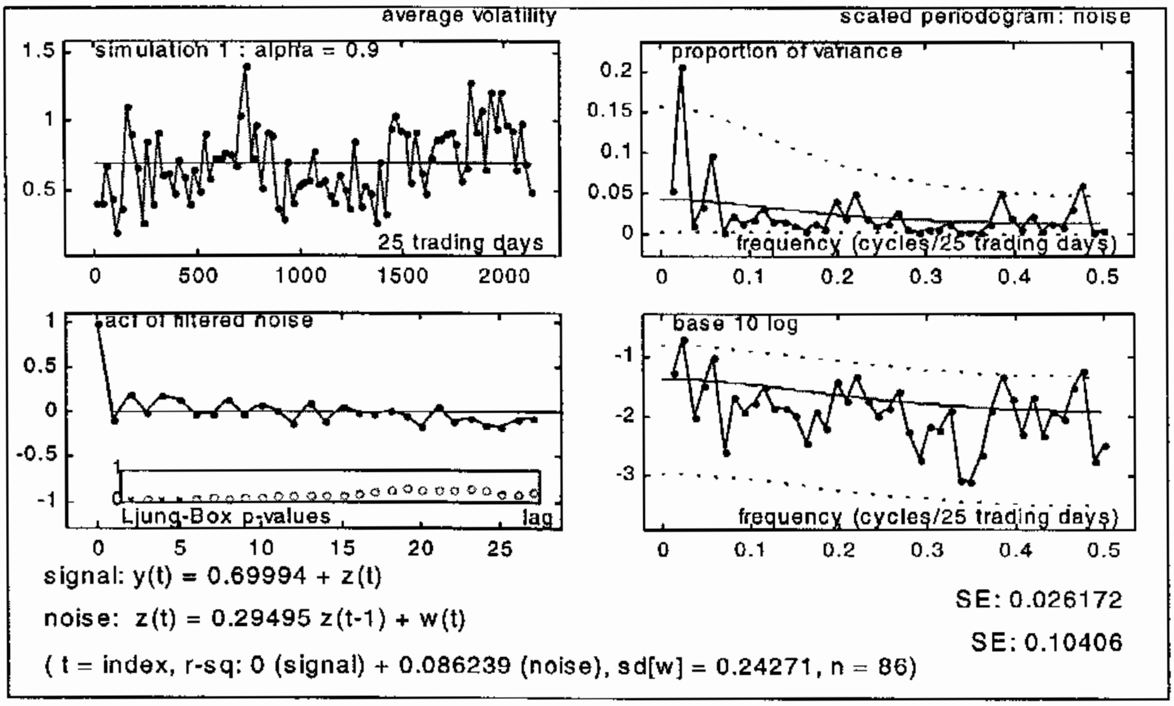

Figure 4.12 shows the time series analysis of the volatility of the first simulation series for $\alpha = 0.95$, for exchange rate returns based on 25-day averages. Again, we test the null hypothesis that the simulation of volatility is constant. The series is fitted well by a first-order autoregression with estimated parameter $a = 0.523$. The standard error of this

estimate is 0.0.093, so the z-statistic for testing the null hypothesis is z = 0.523/0.093 = 5.62, giving a very small p-value (0).

Figure 4.12: Time series analysis of 25-day volatility for simulated returns with $\alpha = 0.95$

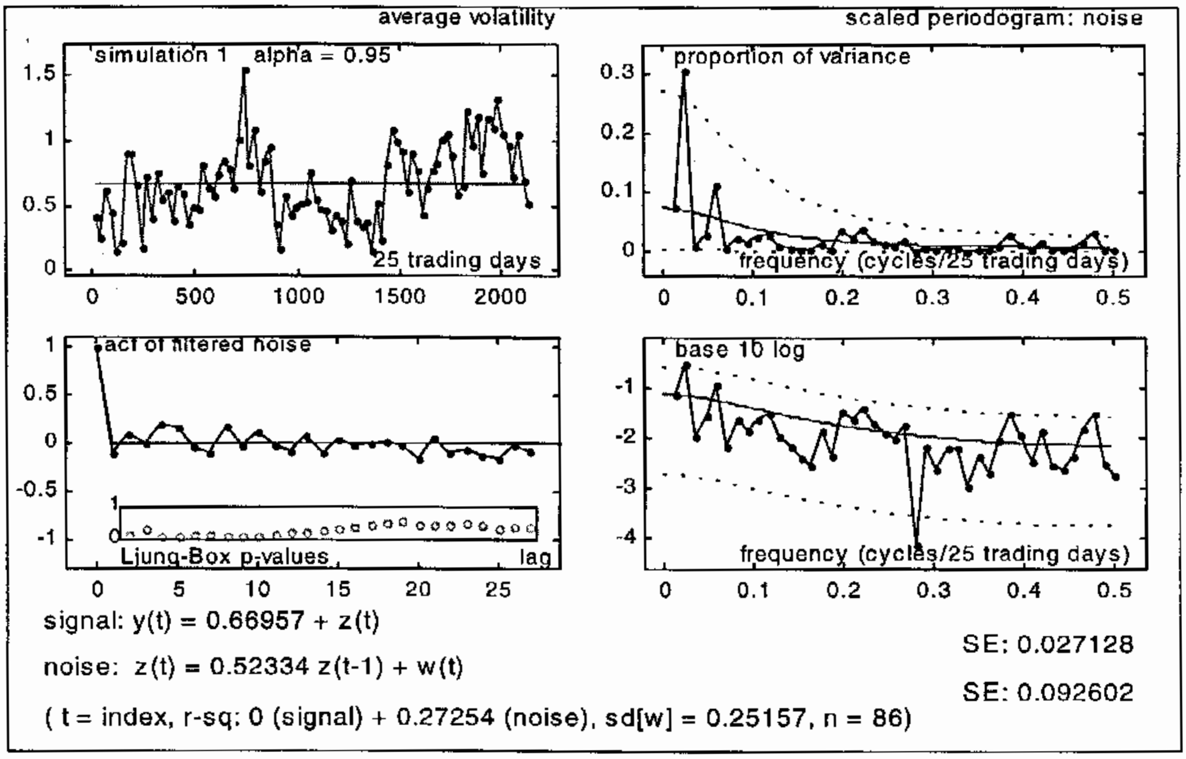

In both Figures 4.8 and 4.9 the small p-values indicate that the volatility is not constant. However, when $\alpha = 0.95$ the fitted autoregressive parameter is substantially larger than those obtained for the data. This suggests that $\alpha = 0.9$ may be the better choice for the parameter.

Summary

We have compared the actual data and the simulated data by using simple model for stochastic volatility.

First we have found that the time series of the volatility for the pound, the yen and the deutsche mark relative to the US dollar are not constant (p-values of 0.0003, 0.021 and 0.0054, respectively).

Second we plot the autocorrelation function of the squared exchange rate returns for pound/dollar, together with theoretical functions obtained $\alpha = 0.95$ and $\alpha = 0.9$, for which the autocorrelations based on the data are mostly positive. The theoretical values are higher than the values obtained for the data for small lags, particularly for $\alpha = 0.95$. This indicates that the limiting model does not provide a satisfactory fit.

Third we estimate the standard deviation and kurtosis coefficients of these simulated exchange rate returns. The volatilities look just like the volatility series based on the data but the estimated skewness coefficients are lees than those of the data.

Finally we have found that in the simulations for the pound, the yen and the deutsche mark relative to the dollar are not constant (p-values of 0.0003, 0.021 and 0.0054, respectively).