

Kaehler, Jürgen

Working Paper — Digitized Version

Modelling and forecasting exchange-rate volatility with ARCH-type models

ZEW Discussion Papers, No. 91-02

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Kaehler, Jürgen (1991) : Modelling and forecasting exchange-rate volatility with ARCH-type models, ZEW Discussion Papers, No. 91-02, Zentrum für Europäische Wirtschaftsforschung (ZEW), Mannheim

This Version is available at:

<https://hdl.handle.net/10419/29445>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

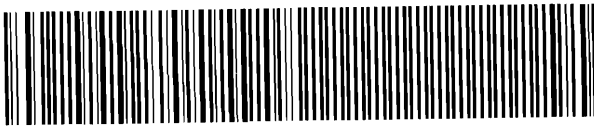
Discussion Paper

Discussion Paper No. 91-02

Modelling and Forecasting Exchange-Rate Volatility with ARCH-Type Models

Juergen Kaehler

W 636 (91.02)



19. DEZ. 1991

Weltwirtschaft
Kiel

W

636-91,2 W

ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

International Finance Series

Modelling and Forecasting Exchange-Rate Volatility with ARCH-Type Models

by

Juergen Kaehler

*Zentrum für Europäische Wirtschaftsforschung (ZEW)
and
Universität Mannheim*

October 1991
(First version: June 1991)

Abstract

The statistical analysis of short-run exchange-rate data shows that there is strong heteroskedasticity and serial dependence of volatility. In addition, the empirical distributions are leptokurtic. The model of generalized autoregressive conditional heteroskedasticity (GARCH) seems to be ideally suited to model these empirical regularities because the model incorporates autocorrelated volatility explicitly and it also implies a leptokurtic distribution. The GARCH model does indeed achieve a reasonably good fit to the exchange-rate data. However, the GARCH model is not able to outperform the naive forecasts of volatility which use the current estimate of the variance from the past data.

Acknowledgements

I wish to thank Volker Marnet, Jörn-Steffen Pischke and the participants of the Econometrics Seminar at the University of Fribourg for helpful comments and suggestions.

1. Introduction

In recent years, an increasing interest in the volatility of financial variables and its implications for the pricing of derivative securities has developed. Nowadays it is so fashionable to talk about "volatility" that this term even begins to replace the familiar terms "standard deviation" and "variance" in the terminology of academics doing research in financial economics. Until quite recently, however, financial economists obtained from modelling volatility as a random variable. In the classic paper of Black and Scholes (1973), for instance, volatility σ (which is a synonym for the instantaneous standard deviation) is assumed to be constant in the differential equation for the price P of the underlying asset

$$(1) \quad dP/P = \mu dt + \sigma dz_1$$

where dz_1 is a Wiener process and μ denotes the instantaneous mean.

Clearly, the assumption of constant volatility is at odds with experience from observing financial markets. However, only in the last few years have financial economists started to develop option pricing models which treat volatility as a stochastic variable (see Hull and White (1987), Johnson and Shanno (1987), Scott (1987), Wiggins (1987), Chesney and Scott (1989), and Melino and Turnbull (1990)). In these models, the diffusion process (1) is supplemented by a diffusion process for the volatility in the form of a geometric generalized Wiener process

$$(2a) \quad d\sigma/\sigma = \phi dt + \gamma dz_2$$

or in the form of an Ornstein-Uhlenbeck process

$$(2b) \quad d\sigma = \theta(\zeta - \sigma)dt + \gamma dz_2$$

However, there are two basic problems with this approach. The first one is analytical. An explicit solution to this model which is independent of risk preferences is only possible if there is either an asset that is instantaneously perfectly correlated with volatility or if volatility is uncorrelated with aggregate consumption. Since both assumptions would in general be regarded as highly unrealistic, the attractiveness of the ingenious Black-Scholes approach to value options through a perfect hedge strategy is lost. Volatility is not a traded asset and since volatility is an unobservable variable, chances are not good that it will ever be traded. Therefore, there is no convincing way to find a perfect hedge for stochastic volatility.

The second drawback of this approach is the fact that the specifications of (2a) and (2b) are entirely ad hoc (as Melino and Turnbull (1990) confess). The specifications are not directed by theoretical considerations or by empirical evidence but rather by analytical convenience. Both (2a) and (2b) are standard stochastic processes which are rather straightforward to work with.

In recent years, there has been a development in the econometrics literature which seems to be well-suited to complement the work on stochastic volatility in the finance literature (see Taylor (1990)). Since the seminal paper of Engle (1982), a rich literature has emerged to model heteroskedasticity (the familiar econometric term for stochastic variance) in a way which bears some resemblance with the univariate ARIMA approach for the mean of a time series. Engle (1982) introduced the autoregressive conditional heteroskedasticity (ARCH) model in which the conditional variance h_t of the variable x_t is a linear function of squared lagged realizations of x_t ¹:

$$(3a) \quad h_t = a_0 + a_1 x_{t-1}^2 + \dots + a_p x_{t-p}^2 \quad \text{with } a_i > 0 \text{ for all } i$$

and

$$(3b) \quad x_t = u_t \sqrt{h_t},$$

where u_t is Gaussian white noise with unit variance.

In applications of the ARCH(p) model, it often turned out that the required lag length p was rather large. This led Bollerslev (1986) to introduce the generalized ARCH model (GARCH for short) whose basic idea is to approximate a long polynomial by a simple rational function, i.e.

$$(4) \quad h_t = a_0 + a_1 x_{t-1}^2 + \dots + a_p x_{t-p}^2 + b_1 h_{t-1} + \dots + b_q h_{t-q}$$

In general, the value of p in (4) will be much smaller than the value of p in (3a).

There are a number of other variants of the ARCH model. Modifying and extending the ARCH model is still a very active research area (see e.g. Nelson (1991), and Harvey, Ruiz and Shephard (1991)). There are also numerous applications of ARCH-type models in finance including the modelling of risk premia and the CAPM with varying covariances.

¹ Throughout this paper I shall assume that the mean of X_t is constant and zero. I shall apply the ARCH-type models only to exchange-rate data and for those series it can be shown that the assumption of a constant mean at zero cannot be rejected (see Kaehler (1989)).

In this paper I explore whether ARCH-type models can be used to overcome the arbitrariness of a volatility specification such as given in (2a) or (2b). Since ARCH-type models are flexible enough to allow a rich dynamic structure, these models may guide the modelling of stochastic volatility. More specifically, I shall examine the forecasting performance of ARCH-type models for foreign exchange-rate series. There are two reasons for looking at forecasting performance. First, an out-of-sample test of a model provides a strong test and is very sensitive to structural instability. Second, from a practical point of view, the main interest of financial management lies in the forecasting of financial variables. For the pricing of foreign-currency options, for example, it is the future volatility of the exchange rate which is the relevant variable.²

I shall only analyse the models with exchange-rate data but the approach taken here is readily extended to other financial variables such as prices of stocks, commodities or bonds. It is a very remarkable fact that there are strong similarities in the statistical properties of different financial variables. This permits to apply the same models to different financial data (see e.g. Taylor (1986)). Before ARCH-type models are estimated and used to forecast volatility of exchange rates, I shall examine statistical properties of four exchange-rate series in some detail in order to explore and quantify the time-series structure and randomness of exchange-rate volatility.

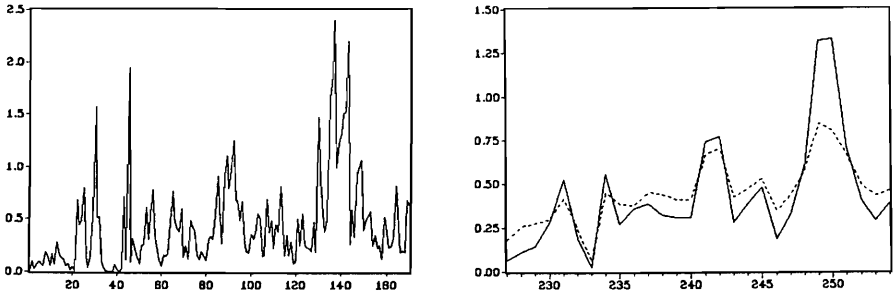
2. Empirical Evidence on Stochastic Volatility

The data to be analysed are the exchange rates of the U.S.dollar against the German mark, the British pound, the Swiss franc and the Japanese yen. The data are on a daily basis but also weekly, monthly and quarterly data are used. In these cases, end-of-period data were derived from daily exchange rates. The data range from July 1st, 1974 to December 31st, 1987. Due to differences in bank holidays between countries, there are different numbers of observations in the daily data: 3386 for the mark, 3417 for the pound, 3392 for the franc and 3365 for the yen. For all currencies, the number of observations in the weekly series is 704, in the monthly series it is 161 and in the quarterly series it is 53. Data source is the IMF's International Financial Statistics, except for the franc, whose exchange rate against the dollar from July 1974 to April 1980 was not published in the International Financial Statistics and was therefore taken from the monthly reports of the Swiss National Bank. The exchange rate dynamics are analysed in the form of $x_t = 100\Delta e_t$ where $\Delta e_t = e_t - e_{t-1}$ and e_t is the logarithm of the exchange rate at time t .

² The results of this study are not directly comparable with those of Taylor (1987) since he analysed future prices instead of spot prices. Furthermore, he used absolute returns and the spread of daily high and low prices to forecast spreads.

In order to provide a first snapshot of the variability of volatility, Figure 1 plots the variances of daily exchange-rate fluctuations x_t in subperiods of length 20 (approximately monthly subperiods) and of length 120 (approximately half-year subperiods) for the pound-dollar rate. The variances show indeed marked variability. However, there does not appear to be a clear pattern in the volatility of subperiods. Periods of turbulence and periods of tranquillity seem to follow one another in a random way. Looking at the display of variances in Figure 1b, one tends to detect a positive trend and cyclical variation in volatility. With only 28 observations, however, this probably reads too much into the data. It should be stressed that the other exchange rates show very similar patterns but they are omitted to save space.

Figure 1
 Variances in subperiods: the pound-dollar rate
 a) monthly subperiods b) half-year subperiods



A more formal test of heteroskedasticity is clearly called for in order to substantiate the randomness of volatility. Here I apply Levene's test which is robust with respect to the underlying distribution of x_t ³. The test is based on a one-way analysis of variance for $w_{kt} = |x_{kt} - \bar{x}_k|$, where \bar{x}_k is the median of the x_t 's in the k -th subsample⁴. The null hypothesis of equal variances in K subsamples, $H_0: \sigma_1^2 = \dots = \sigma_K^2$, will be rejected if the test statistic Λ exceeds the $(1 - \alpha)$ -quantile of the F-distribution with $K-1$ and $T-K$ degrees of freedom, where T is the number of observations. It is difficult, however, to test H_0 in a rigorous way without an a-priori perception of the number and

³ As I will show later, the assumption of a normal distribution, on which most parametric tests are based, is quite questionable for short-run exchange-rate dynamics.

⁴ In Figure 1b, w_{kt} is plotted as a dashed line.

the size of subsamples. Since there is no natural division of observations into K subgroups, different divisions will be employed. The data will be subdivided into sequences of equal length with length of 20, 60, 120 and 240, respectively. This corresponds roughly to time intervals of a month, a quarter, half a year and a year.

The results are reported in Table 1. For all entries of Table 1 the Λ -estimates fall far into the upper tail of the corresponding F-distribution. In fact, for all 16 Λ -values of this table, the empirical significance level is at least of order 10^{-11} . However, apart from indicating that there is extremely strong evidence for heteroskedasticity in the data, the Levene test does not identify any structure for heteroskedasticity, nor does the test identify the subsamples with abnormal variances.

Table 1
Levene's test for homogeneity of variance

	mark	pound	franc	yen
month	5.2	5.7	6.1	6.5
quarter	11.0	13.3	10.4	13.3
half-year	15.6	19.0	14.1	22.9
year	25.8	32.7	21.8	42.8

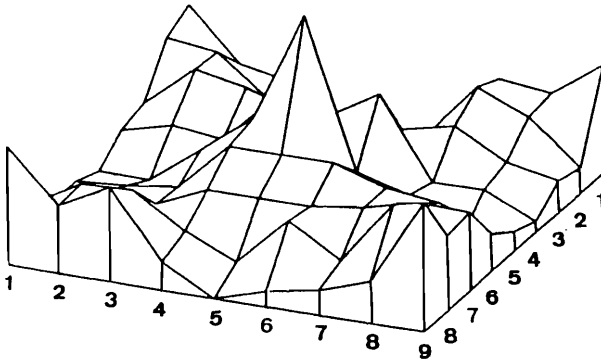
On the other hand, it is often claimed that there is positive correlation of volatility in financial markets. As Mandelbrot (1963, p. 418) put it in his seminal paper: "Large (price) changes tend to be followed by large (price) changes - of either sign - and small (price) changes tend to be followed by small (price) changes". This assertion can be examined in a direct way within a simple Markov-chain model.

Let the observations x_t be classified in an ascending order into J quantiles where the first quantile contains the largest depreciations of the dollar and the J -th quantile contains the largest appreciations. The quantiles are chosen such that all quantiles have the same number of observations. One may then count the number of times that an observation x_t falls into quantile J_i and x_{t+1} falls into quantile J_j and denote this number by n_{ij} . If the x_t are independent and identically distributed then the expected value is $\tilde{n}_{ij} = (T-1)/J^2$.

A typical empirical transition matrix of n_{ij} 's is displayed in Figure 2. It shows the data of daily changes in the pound-dollar rate classified into 9 quantiles. The height of the three-dimensional body is proportional to n_{ij} . There are 5 main peaks; a dominant one with $n_{55} = 109$ and 4 peaks in the corners with $n_{11} = 65$, $n_{19} = 66$, $n_{91} = 70$ and $n_{99} = 74$. There is also a side peak with $n_{46} = 69$. For all entries, the expected number is $\tilde{n}_{ij} = 42.2$. The main peaks can be interpreted in terms of periods of tranquillity and

is $\bar{n}_{ij} = 42.2$. The main peaks can be interpreted in terms of periods of tranquillity and periods of turbulence. The dominant peak n_{55} gives the number of cases where a small $|x_t|$ is followed by another small $|x_{t+1}|$. Likewise, n_{11} is the number of cases where a strong depreciation of the dollar against the pound was followed by another strong depreciation and n_{99} is the number of pairs of strong appreciation. On the other hand, n_{19} cases could be counted where a strong appreciation followed a strong depreciation et vice versa for n_{91} .

Figure 2
Markov transition matrix: daily pound dollar rate



This indicates that in turbulent periods there can be a strong reaction in the foreign-exchange market of either sign, i.e. a strong exchange-rate movement in period t lowers the probability of small or moderate movements in $t+1$ and increases the probabilities both for a strong depreciation and a strong appreciation, confirming Mandelbrot's observation.

A rigorous test of this observation is provided by a familiar Chi-squared goodness-of-fit test applied to the empirical transition matrices. The test results for all exchange-rate series are reported in Table 2.

For daily, weekly, and monthly series, J is equal to 5 and for quarterly series J is equal to 3. Obviously, there is strong rejection of independence for daily and weekly data but only very weak evidence against independence for longer-term exchange-rate fluctuations. The choice of J does not seem to have an influence upon this result.

Performing the same tests with $J = 9$ for daily and $J = 7$ for weekly data resulted in the same rejections of independence at very high significance levels. In fact, increasing J brought an increase in all test statistics.

Table 2
Testing for independence in Markov chains

	mark	pound	franc	yen
day	128.8 ***	316.0 ***	100.6 ***	383.0 ***
week	42.4 ***	66.9 ***	41.4 ***	90.5 ***
month	13.5	18.9	15.4	21.3
quarter	8.8 *	5.7	4.3	12.2 **

Significance levels: * 10 percent, ** 5 percent, *** 1 percent.

The advantage of testing the transition matrix in this way for first-order serial independence is that no strong distributional assumptions about X_t are required. The test, however, has two drawbacks. First, classifying exchange-rate data into quantiles wastes a lot of information and, second, higher-order serial dependence is not examined. The phenomenon that small fluctuations tend to be followed by small fluctuations and large fluctuations by large ones of either sign can be measured without loss of information by estimating the autocorrelation function (ACF) of squared observations x_t^2 . In addition, the ACF, which is familiar from time-series analysis, is able to detect a rich pattern of dependence in variances.

Table 3 reports results from applying the Ljung-Box test to autocorrelations up to order 15. The test statistic $Q(15)$ for all four exchange rates at four different time horizons each is reported as the upper number in Table 3. It is evident that there is strong rejection of the H_0 of no serial dependence in variances for daily and weekly data only. In the daily series, the ACF for squared exchange-rate movements is significant at all lags up to 15 for all four exchange rates. For weekly data, the estimated autocorrelation coefficients exceed the conventional confidence limits of $\pm 2\sqrt{T}$ at various lags. The number of significant autocorrelations and (after the slash) partial autocorrelations is given below the corresponding Q statistics.

According to Bollerslev (1988), the ACF and the partial ACF for squared data can be used in the same way as in conventional ARIMA models to identify the order of the autoregressive (AR) component and of the moving-average (MA) component in corresponding models for variances. The great number of significant autocorrelation and partial autocorrelation coefficients suggests that there is some structure in the time pattern of variance which is worth modelling and that a mixed AR and MA process might be adequate.

Table 3
Results from the ACF of squared data

	mark	pound	franc	yen
day	355.2 *** 15/6	507.3 *** 15/9	516.0 *** 15/10	432.2 *** 15/7
week	61.5 *** 5/4	123.9 *** 7/7	98.2 ** 8/3	52.8 *** 6/3
month	12.3 1/1	12.9 1/1	11.8 1/0	25.0 ** 2/1
quarter	12.0 0/0	7.4 0/0	12.1 0/0	5.4 0/0

Significance levels: see Table 2.

To sum up, the statistical analysis of volatility reveals that there is strong heteroskedasticity in the data and that volatility is positively autocorrelated up to large lags. These properties, however, apply to short-run exchange-rate dynamics only. There is no strong evidence for random volatility in monthly or quarterly data.

Another strong empirical property, which is common to speculative prices, is the non-normality of short-run price dynamics (see e. g. Taylor (1986)). It has repeatedly been found that these variables have excess kurtosis, i. e. kurtosis which is significantly greater than 3 (the value for a normal distribution). Since kurtosis β_2 , defined as the fourth central moment divided by the square of the variance, measures both tail-weight and peakedness, excess kurtosis indicates excessive mass in the tails or at the centre of the empirical distribution. A test of $H_0: \beta_2 = 3$ is a test of mesokurtosis with the two-sided alternatives of platykurtic ($\beta_2 < 3$) and leptokurtic ($\beta_2 > 3$) distributions. The values of β_2 are reported in Table 4 for each series of x_t and test statistics, which have an approximate normal distribution under H_0 , are given below in brackets. As the table shows, there is extremely strong leptokurtosis in the daily and weekly series. In the monthly series, the H_0 of mesokurtosis can be rejected at the 0.05 level for 3 exchange rates, whereas no rejection of H_0 is possible for any of the quarterly series. This means that leptokurtosis is essentially a property of short-run exchange-rate dynamics. It is only moderately inherent in monthly series and vanishes completely in quarterly data.

Table 4
Test for mesokurtosis

	mark	pound	franc	yen
day	8.32 *** (19.93)	8.36 *** (20.07)	8.89 *** (20.70)	8.00 *** (19.93)
week	5.84 *** (7.21)	7.36 *** (8.73)	4.96 *** (5.93)	7.03 *** (8.45)
month	3.87 ** (2.01)	4.15 *** (2.39)	4.19 *** (2.45)	3.62 (1.61)
quarter	2.67 (-0.22)	2.72 (-0.11)	2.77 (-0.02)	2.62 (-0.33)

Significance levels: see Table 2.

The results of the statistical analysis can succinctly be summarized as follows: Short-run exchange-rate dynamics (i.e. daily and weekly changes) are characterized by a time pattern in heteroskedasticity as well as peakedness and fat tails in distribution whereas medium-run exchange-rate dynamics (i.e. monthly and quarterly changes) show no heteroskedasticity or serial dependence and have a frequency distribution which is approximately normal.

The relevant time span for the modelling of financial volatility along the lines of (2a) or (2b) would be the very short run since dz_2 is a continuous-time variable and σ is the instantaneous standard deviation. A discrete-time approximation of (2) should, therefore, use a time-interval as short as possible. Thus, the relevant statistical properties should be those of the short-run exchange-rate dynamics. In the next section, a number of models will be considered which are able to capture some or all of these statistical properties.

3. Modelling Stochastic Volatility

The early contributions to the stochastic modelling of price dynamics in financial markets aimed to capture leptokurtosis in their model. This research was initiated by Mandelbrot (1963) and Fama (1965). Mandelbrot (1963) suggested to apply the family of stable distributions which leads to a very general model since this family is derived from a generalization of the central limit theorem. If one drops the assumption of finite variance in the conditions of the central limit theorem, one gets stable Paretian distributions as the only possible limit distributions for sums of independent and identically

distributed random variables. The drawback of stable Paretian distributions, however, is their lack of closed forms for the density function (with only a few exceptions) and hence their analytical awkwardness.

The family of stable distributions was later supplemented by other probability models which share the property that they can be described as compound distributions.⁵ They all assume that the variable X_t follows a normal distribution with a stochastic variance. The mixture of normal distributions assumes that the variance has a Bernoulli distribution (see Ball and Torous (1983)), Student's distribution attaches an inverted gamma distribution to the variance (see Praetz (1972)), Clark (1973) assumes that the variance has a log-normal distribution and the model of Press (1967) can be reformulated in a way such that it implies a Poisson distribution for the variance. Even the class of symmetric stable distributions can be derived as a compound distribution where the variance follows a sub-class of stable distributions.

The application of these compound distributions can be motivated by the fact that they imply leptokurtosis⁶ and are, therefore, consistent with a strong and robust empirical property of speculative price. Although these models specify the variance as a random variable, they are not able to explain heteroskedasticity or serial dependence in variances. On the other hand, serial dependence in variances is explicitly incorporated in the ARCH(p) and GARCH(p, q) models of (3) and (4). Furthermore, Milhoj (1985) showed for the ARCH(p) model and Bollerslev (1986) showed for the GARCH(1,1) model that these models imply leptokurtosis of X_t . Thus, the ARCH-type models are consistent with the major stylized facts of short-run price dynamics in financial markets and appear to be ideal candidates to model these price dynamics.

A priori nothing can be said about the number of lags p which need to be included into the specification of (3). The results from the empirical autocorrelation functions for squared data, however, indicate that high order dependence is present in short-run dynamics. Therefore, it is best to choose the lag-length with a model selection criterion and I applied the Schwarz information criterion (SIC), defined by $SIC = r \log T - 2L^*$ where r is the number of parameters and L^* is the logarithm of the maximised likelihood.

I estimated ARCH models up to order $p=25$ and identified models with lags between 11 and 20 for daily data and with lags between 3 and 12 for weekly data. This suggests to find a more parsimonious parametrization of the model. There are basically two possibilities. First, one can impose a restriction on the a_i 's in the form of linearly or geometrically declining weights (or some other functional form) as suggested by Engle (1982). Hsieh (1989), however, found that the restrictions of linearly and of

⁵ Clark (1973) calls these kind of models "subordinated stochastic processes".

⁶ This general result follows from Beale and Mallows (1959) in which they showed that scale mixtures of normal distributions are leptokurtic.

geometrically declining weights were rejected by likelihood ratio (LR) tests for daily exchange-rate data. The alternative is to find a more parsimonious parametrization by applying the GARCH (p, q) model. I chose the simple specification $p=1$ and $q=1$.

The choice between the ARCH(p), where p is determined by SIC, and the GARCH(1,1) model can again be made by applying information criteria. Table 5 reports the comparison between ARCH and GARCH for the exchange-rate data by SIC ⁷. Only in one out of 8 cases is SIC lower, and hence better, for the ARCH model than for the GARCH model and this result would clearly favour the GARCH model. In the following, I will therefore only report results from the GARCH model.

Table 5
Comparison between ARCH(p) and GARCH(1,1) models by SIC

		mark	pound	franc	yen
day	ARCH	6097.5 (11)	6030.6 (20)	7323.9 (12)	5427.1 (11)
	GARCH	6064.3	5981.2	7299.7	5430.1
week	ARCH	2463.5 (3)	2454.5 (4)	2701.6 (6)	2300.3 (12)
	GARCH	2442.4	2422.7	2663.0	2237.4

Estimation of the GARCH(1,1) model by Maximum Likelihood methods is quite straightforward. The estimates are shown in Table 6 and standard errors are given in brackets. The parameters a_1 and b_1 can be estimated with quite high precision, especially in daily data. In most cases, a_1 is close to 0.1 and b_1 is close to 0.9. Since the mean lag of conditional variance effects is given by $(1 - b_1)^{-1}$, the high value of b_1 implies that there is strong persistence in variances.

The fact that the sum of a_1 and b_1 is close to 1 leads to the issue of stationarity. Bollerslev (1986) showed that a GARCH(p, q) process is second-order stationary if and only if

$$(5) \quad a_1 + \dots + a_p + b_1 + \dots + b_q < 1.$$

According to Table 6, there are several series for which $a_1 + b_1 > 1$. This violation of the stationarity condition has been observed repeatedly in applications of the GARCH

⁷The lag p of the ARCH models is given in brackets.

Table 6
Estimates of the GARCH(1,1) model

		mark	pound	franc	yen
day	a_0	0.656 (0.128)	0.675 (0.109)	0.856 (0.179)	0.766 (0.120)
	a_1	0.169 (0.015)	0.135 (0.013)	0.143 (0.013)	0.190 (0.021)
	b_1	0.833 (0.013)	0.864 (0.011)	0.857 (0.012)	0.820 (0.017)
	LR	970.3 ***	908.5 ***	1015.1 ***	901.6 ***
	β_2	4.80 (11.92)	10.58 (22.63)	6.29 (16.31)	17.44 (26.87)
	Q(15)	22.86 *	8.43	54.08 ***	3.19
week	a_0	1.71 (1.34)	4.16 (1.68)	1.83 (1.22)	0.52 (0.38)
	a_1	0.095 (0.027)	0.114 (0.029)	0.095 (0.017)	0.076 (0.013)
	b_1	0.906 (0.027)	0.878 (0.025)	0.907 (0.016)	0.932 (0.011)
	LR	111.4 ***	103.0 ***	119.3 ***	133.8 ***
	β_2	4.29 (4.63)	7.03 (8.45)	4.26 (4.56)	9.81 (10.31)
	Q(15)	29.56 **	8.03	11.66	4.43

Significance levels: see Table 2. The values of α_0 and their standard errors are multiplied by 100.

model to financial data. This lead Engle and Bollerslev (1986) to extend the GARCH model to the case where variances are non-stationary. The integrated GARCH model, IGARCH for short, obtains if the polynomial equation

$$(6) \quad 1 - a_1 v - \dots - a_p v^p - b_1 v - \dots - b_q v^q = 0$$

has at least one unit root, whereas the GARCH model requires that all roots lie outside the unit circle of the complex plane. I do not want to pursue the idea of integration in variance here further because the statistical properties of the IGARCH model are not yet fully developed (see also the discussion of the Engle-Bollerslev paper in volume 5 of *Econometric Reviews*).

I also performed a LR test for the GARCH(1,1) model against the H_0 of a simple normal distribution, i.e. the H_0 implies that $p = 0$ and $q = 0$. The LR test rejects Gaussian white noise very strongly against the GARCH(1,1) model for all series (see Table 6). This result, of course, casts serious doubt on the appropriateness of the diffusion process in (1) which is part of the Black-Scholes model and of the standard currency-options model (see e.g. Garman and Kohlhagen (1983)).

It is also instructive to analyse the "residuals" $\hat{u}_t = x_t / \hat{h}_t^{1/2}$ in order to see whether the GARCH(1,1) model fits the data. According to the model, u_t has a standard normal distribution. Table 6 reports the kurtosis of the residuals and in brackets the value of the test statistic for a test of mesokurtosis. The H_0 of mesokurtosis can be rejected at

very high significance levels for all series. Even more surprisingly, leptokurtosis increases substantially in \hat{u}_t , as compared with x_t , for both yen series and the daily pound series (cf Table 4). For the daily yen series, kurtosis more than doubles. This points to a weakness of this model which might be due to a wrong distributional assumption about u_t , or other misspecifications.

Table 6 also reports the Ljung-Box statistic Q for squared residuals at lag 15. As compared with the same statistics for the raw data (see Table 3), there is a dramatic drop in Q for the residuals of the GARCH(1,1) model. In all cases, the Q for the residuals is less than 10 percent of the Q for the raw data. Only for the daily franc series is the Q of the residuals significant at the 1 percent level. One may conclude, therefore, that the GARCH(1,1) model captures the empirical volatility dynamics well.

Figure 3

Squared data and conditional variance: weekly pound-dollar rate

a) Squared data

b) Conditional variance

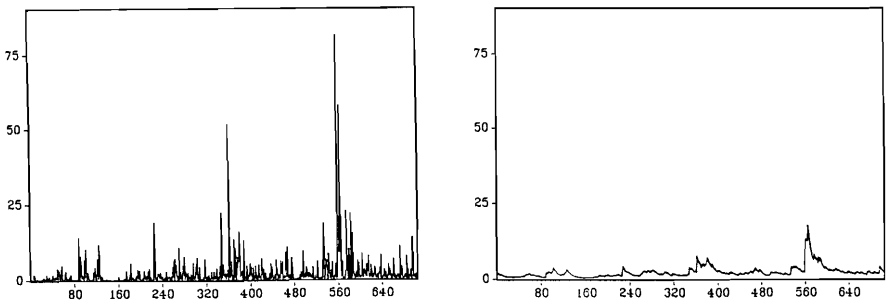


Figure 3 displays the series x_t^2 and \hat{h}_t for the weekly pound series in order to give a more detailed picture of the fit the GARCH model achieves. The comparison between the two series shows that their patterns are very similar but the amplitude of \hat{h}_t is much smaller than the amplitude of x_t^2 . This explains why the kurtosis of the residuals is very high.

4. Forecasting Volatility

From the point of view of financial management, there is, for obvious reasons, a special interest in the forecasting of financial-market prices. Forecasting experiments with the estimated GARCH models can, in addition, shed some light on the structural stability of the models. The compound distribution models referred to in the last-section do not

imply any serial correlation of volatility. Hence, these models would predict a constant volatility and the best predictor from these models would simply be the volatility from past realizations. These models serve as a benchmark for the GARCH model. On the other hand, the GARCH model provides non-trivial forecasts of future variances.

I applied the same strategy for measuring forecasting performance as was used by Meese and Rogoff (1983) in their comparison with respect to the forecastability of the mean. Thus, I estimated the models, on a "rolling basis" as in Meese and Rogoff (1983). For the daily data, the GARCH model was first estimated for the observations from $t = 1$ to $T = 1000$. Forecasts were made for the next 20 time periods and the forecasts were compared with x_{T+i}^2 . In the next step, 100 observations were added, parameters were re-estimated and forecasts were again compared with observations. In this way, parameters and forecasts were computed 23 times for each daily series. For weekly data, the first estimation period includes observations up to $T = 220$ and on each step 20 observations were added to the previous subsample. The forecast horizon includes each of the next 20 weeks. This gives 24 forecast experiments for each of the weekly series.

The forecasts of the GARCH model and the "naive" models (constant variance) are compared with respect to mean errors and with respect to root mean square errors (RMSE). The results are summarized in Tables 7 and 8. Note that the mean errors and RMSE's are averaged over all forecast horizons.

Table 7
Comparison of models by forecasting variances: mean errors

		mark	pound	franc	yen
day	GARCH	0.116	0.165	0.464	0.144
	$\sigma^2(t)$	-0.217	-0.184	-0.005	-0.269
week	GARCH	0.077	2.571	0.985	0.188
	$\sigma^2(t)$	-1.193	-1.060	-0.605	-0.765

Table 8
Comparison of models by forecasting variances: RMSE

		mark	pound	franc	yen
day	GARCH	1.32	1.16	2.16	1.35
	$\sigma^2(t)$	1.20	1.02	1.27	1.20
week	GARCH	5.09	7.63	6.08	3.87
	$\sigma^2(t)$	4.75	4.95	5.60	3.66

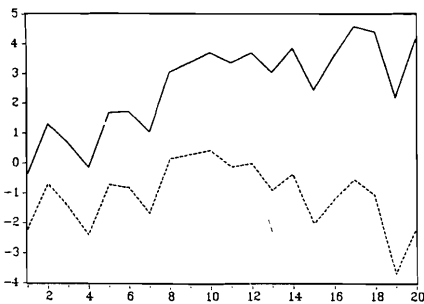
As regards bias of the forecasts, the results are mixed. For daily data, the GARCH model achieves in 3 cases a lower absolute mean error but for weekly data it is not clear which forecast function is better. The most remarkable result, however, is that the "naive" forecasts of variances are clearly better than GARCH forecasts in terms of precision. For all daily and weekly series, the "naive" models perform better than the GARCH models according to the RMSE criterion, sometimes quite substantially so. In a way, these results reproduce Meese and Rogoff's results for the forecasting of means. Here it is shown that the non-forecastability extends to variances. It is also interesting to note that the superiority of the random walk model over asset market models in forecasting the mean is more obvious with respect to the average (over all 20 forecast horizons) RMSE than with respect to the average mean errors (see Meese and Rogoff (1983)). Thus, there is another correspondence between their results and the results on variances presented here.

In order to gain more insight into the forecasting performance, Figure 4 plots mean errors and RMSE at forecast horizons 1 to 20 for the weekly pound series. As regards bias, the mean errors of the GARCH model are generally positive whereas they tend to be negative for the naive model. In the case of the weekly pound series, GARCH mean errors are only negative for 1-step and 4-step ahead forecasts while the mean errors of the naive model are only positive for forecast horizons from 8 to 10. However, the pattern of mean forecast errors in Figure 4a is very similar for both models.

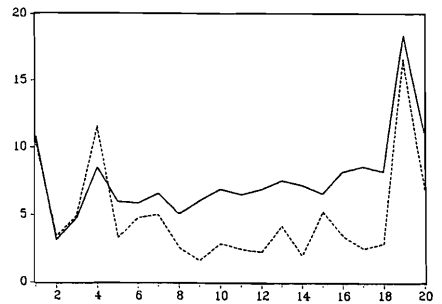
Figure 4

Forecast errors of variances at different time horizons: weekly pound-dollar rate

a) mean errors



b) RMSE



As regards precision of the forecasts, the RMSE of the GARCH model tends to be higher than the RMSE of the constant-variance model in Figure 4b. At forecast horizons 1 to 3, both models have virtually the same RMSE and at forecast horizon 4, the GARCH model performs better than the naive model. However, the naive model is better at all other forecast horizons. When all daily and weekly exchange-rate series

are taken together, the constant-variance model is on average, i.e. over all forecast horizons, better than the GARCH model as regards RMSE. However, constant variance forecast are not better than GARCH forecasts when volatility is high.

5. Concluding Remarks

In this paper I addressed the question of how to model and forecast the volatility of financial prices. The application was only to exchange-rate data but it should be emphasized that the methods of analysis are readily extended to other financial-market data. Since there is great similarity between statistical properties of different speculative prices, one may expect to find results for these other prices which are very similar to the results reported here.

The analysis of exchange-rate data revealed some interesting statistical properties of short-run dynamics. There is strong heteroskedasticity and serial dependence of volatility. In addition, there are more very large and very small exchange-rate movements than expected under a normal distribution, i.e. the empirical distributions are leptokurtic. The GARCH model seems to be ideally suited to model these data, because it incorporates autocorrelated volatility explicitly and it also implies a leptokurtic distribution.

The GARCH model does indeed achieve a reasonably good fit to the data since it captures serial dependence of volatilities. However, the results from the forecast experiment are striking. The GARCH model is not able to outperform the naive forecast which uses the current estimate of the variance from the past data. Statistically, this result is related to the fact, that there seems to be a unit root in the variance of the data (see equation (6)), i.e. not only the mean of the exchange-rate level seems to be on a random walk but also the variance of exchange-rate dynamics.

From an econometric point of view, the poor forecasting performance is a great disappointment. But from the point of view of financial management this result is actually good news. It implies that financial analysts should not worry too much about stochastic volatilities. The current practice to estimate volatility by the historical standard deviation is obviously not inferior to other, more refined approaches based on new econometric techniques.

References

- Ball, C.A. and W.N. Torous (1983): "A Simplified Jump Process for Common Stock Returns", *Journal of Financial and Quantitative Analysis*, 18, 53-61.
- Beale, E.M. and C.L. Mallows (1959): "Scale Mixing of Symmetric Distributions with Zero Means", *Annals of Mathematical Statistics*, 30, 1145-1151.
- Black, F. and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 81, 637-654.
- Bollerslev, T. (1986): "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T. (1988): "On the Correlation Structure for the Generalized Autoregressive Conditional Heteroskedastic Process", *Journal of Time Series Analysis*, 9, 121-131.
- Chesney, M. and L. Scott (1989): "Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model", *Journal of Financial and Quantitative Analysis*, 24, 267-284.
- Clark, P.K. (1973): "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices", *Econometrica*, 41, 135-155.
- Englè, R.F. (1982): "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica*, 50, 987-1008.
- Engle, R.F. and T. Bollerslev (1986): "Modelling the Persistence of Conditional Variances", *Econometric Reviews*, 5, 1-50.
- Fama, E.F. (1965): "The Behavior of Stock-Market Prices", *Journal of Business*, 38, 34-105.
- Garman, M.B. and S.W. Kohlhagen (1983): "Foreign Currency Option Values", *Journal of International Money and Finance*, 2, 231-237.
- Harvey, A., E. Ruiz and N. Shepard (1991): "Modelling Volatility: Some Alternatives to ARCH", mimeo, London School of Economics, Statistical Department.
- Hsieh, D.A. (1989): "Modeling Heteroscedasticity in Daily Foreign-Exchange Rates", *Journal of Business and Economic Statistics*, 7, 307-317.
- Hull, J. and A. White (1987): "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance*, 42, 281-300.
- Johnson, H. and D. Shanno (1987): "Option Pricing when the Variance is Changing", *Journal of Financial and Quantitative Analysis*, 22, 143-151.
- Kaehler, J. (1989): "Statistical Properties of Exchange Rates", Universität Mannheim, Institut für Volkswirtschaftslehre und Statistik, Discussion Paper No. 286-89.
- Mandelbrot, B. (1963): "The Variation of Certain Speculative Prices", *Journal of Business*, 36, 394-419.
- Meese, R. A. and K. Rogoff (1983): "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample ?", *Journal of International Economics*, 14, 3-24.
- Melino, A. and S. M. Turnbull (1990): "Pricing Foreign Currency Options with Stochastic Volatility", *Journal of Econometrics*, 45, 239-265.
- Milhoj, A. (1985): "The Moment Structure of ARCH Processes", *Scandinavian Journal of Statistics*, 12, 281-292.
- Nelson, D. B. (1991): "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica*, 59, 347-370.
- Praetz, P.D. (1972): "The Distribution of Share Price Changes", *Journal of Business*, 40, 317-335.
- Press, S.J. (1967): "A Compound Events Model for Security Prices", *Journal of Business*, 40, 317-335.
- Scott, L. O. (1987): "Option Pricing when the Variance Changes Randomly: Theory, Estimation and an Application", *Journal of Financial and Quantitative Analysis*, 22, 419-438.

- Taylor, S. J. (1986): "Modelling Financial Time Series". Chichester: Wiley.
- Taylor, S. J. (1987): "Forecasting the Volatility of Currency Exchange Rates", *International Journal of Forecasting*, 3, 159-170.
- Taylor, S. J. (1990): "Modelling Stochastic Volatility", mimeo, University of Lancaster, Department of Accounting and Finance.
- Wiggins, J. B. (1987): "Option Values under Stochastic Volatility: Theory and Empirical Estimates", *Journal of Financial Economics*, 19, 351-372.